

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)[End of Result Set](#) [Generate Collection](#) [Print](#)

L6: Entry 7 of 7

File: USPT

Aug 6, 2002

DOCUMENT-IDENTIFIER: US 6430539 B1

TITLE: Predictive modeling of consumer financial behaviorAbstract Text (1):

Predictive modeling of consumer financial behavior is provided by application of consumer transaction data to predictive models associated with merchant segments. Merchant segments are derived from consumer transaction data based on co-occurrences of merchants in sequences of transactions. Merchant vectors representing specific merchants are clustered to form merchant segments in a vector space as a function of the degree to which merchants co-occur more or less frequently than expected. Each merchant segment is trained using consumer transaction data in selected past time periods to predict spending in subsequent time periods for a consumer based on previous spending by the consumer. Consumer profiles describe summary statistics of consumer spending in and across merchant segments. Analysis of consumers associated with a segment identifies selected consumers according to predicted spending in the segment or other criteria, and the targeting of promotional offers specific to the segment and its merchants.

Brief Summary Text (3):

The present invention relates generally to analysis of consumer financial behavior, and more particularly to analyzing historical consumer financial behavior to accurately predict future spending behavior, and more particularly, future spending in specifically identified data-driven industry segments.

Brief Summary Text (6):

Conventional means of determining consumer interests have generally relied on collecting demographic information about consumers, such as income, age, place of residence, occupation, and so forth, and associating various demographic categories with various categories of interests and merchants. Interest information may be collected from surveys, publication subscription lists, product warranty cards, and myriad other sources. Complex data processing is then applied to the source of data resulting in some demographic and interest description of each of a number of consumers.

Brief Summary Text (7):

This approach to understanding consumer behavior often misses the mark. The ultimate goal of this type of approach, whether acknowledged or not, is to predict consumer spending in the future. The assumption is that consumers will spend money on their interests, as expressed by things like their subscription lists and their demographics. Yet, the data on which the determination of interests is made is typically only indirectly related to the actual spending patterns of the consumer. For example, most publications have developed demographic models of their readership, and offer their subscription lists for sale to others interested in the particular demographics of the publication's readers. But subscription to a particular publication is a relatively poor indicator of what the consumer's spending patterns will be in the future.

Brief Summary Text (10):

Yet another problem with conventional approaches is that categorization of purchases is often based on standardized industry classifications of merchants and business, such as the SIC codes. This set of classification is entirely arbitrary, and has little to do with actual consumer behavior. Consumers do not decide which merchants to purchase from based on their SIC code. Thus, the use of arbitrary classifications to predict financial behavior is doomed to failure, since the classifications have little meaning in the actual data of consumer spending.

Brief Summary Text (12):

Accordingly, what is needed is the ability to model consumer financial behavior based on actual historical spending patterns that reflect the time-related nature of each consumer's purchase. Further, it is desirable to extract meaningful classifications of merchants based on the actual spending patterns, and from the combination of these, predict future spending of an individual consumer in specific, meaningful merchant groupings.

Brief Summary Text (13):

In the application domain of information, and particularly text retrieval, vector based representations of documents and words is known. Vector space representations of documents are described in U.S. Pat. No. 5,619,709 issued to Caid et. al, and in U.S. Pat. No. 5,325,298 issued to Gallant. Generally, vectors are used to represent words or documents. The relationships between words and between documents is learned and encoded in the vectors by a learning law. However, because these uses of vector space representations, including the context vectors of Caid, are designed for primarily for information retrieval, they are not effective for predictive analysis of behavior when applied to documents such as credit card statements and the like. When the techniques of Caid were applied to the prediction problems, it had numerous shortcomings. First, it had problems dealing with high transaction count merchants. These are merchants whose names appear very frequently in the collections of transaction statements. Because Caid's system downplays the significance of frequently appearing terms, these high transaction frequency merchants were not being accurately represented. Excluding high transaction frequency merchants from the data set however undermines the system's ability to predict transactions in these important merchants. Second, it was discovered that past two iterations of training, Caid's system performance declined, instead of converging. This indicates that the learning law is learning information that is only coincidental to transaction prediction, instead of information that is specifically for transaction prediction. Accordingly, it is desirable to provide a new methodology for learning the relationships between merchants and consumers so as to properly reflect the significance of the frequency with which merchants appears in the transaction data.

Brief Summary Text (15):

The present invention overcomes the limitations of conventional approaches to consumer analysis by providing a system and method of analyzing and predicting consumer financial behavior that uses historical, and time-sensitive, spending patterns of individual consumers to create both meaningful groupings (segments) of merchants which accurately reflect underlying consumer interests, and a predictive model of consumer spending patterns for each of the merchant segment. Current spending data of an individual consumer or groups of consumers can then be applied to the predictive models to predict future spending of the consumers in each of the merchant clusters.

Brief Summary Text (16):

In one aspect, the present invention includes the creation of data-driven grouping of merchants, based essentially on the actual spending patterns of a group of consumers. Spending data of each consumer is obtained, which describes the spending patterns of the consumers in a time-related fashion. For example, credit card data demonstrates not merely the merchants and amounts spent, but also the sequence in which purchases were made. One of the features of the invention is its ability to use the co-occurrence of purchases at different merchants to group merchants into meaningful merchant segments. That is, merchants which are frequently shopped at within some number of transactions or time period of each other reflect a meaningful cluster. This data-driven clustering of merchants more accurately describes the interests or preferences of consumers.

Brief Summary Text (17):

In a preferred embodiment, the analysis of consumer spending uses spending data, such as credit card statements, and processes that data to identify co-occurrences of purchases within defined co-occurrence windows, which may be based on either a number of transactions, a time interval, or other sequence related criteria. Each merchant is associated with vector representation; the initial vectors for all of the merchants are randomized to present a quasi-orthogonal set of vectors in a merchant vector space. Each consumer's transaction data reflecting their purchases (e.g. credit card statements, bank statements, and the like) is chronologically organized to reflect the general order in which purchases were made at the merchants. Analysis of each consumer's transaction data in various co-occurrence windows identifies which merchants co-occur. For each pair of merchants, their respective merchant vectors are updated in the vector

space as a function of their frequency of their co-occurrence. After processing of the spending data, the merchant vectors of merchants which are frequented together are generally aligned in the same direction in the merchant vector space. Clustering techniques are then applied to find clusters of merchants based on their merchant vectors. These clusters form the merchant segments, with each merchant segment having a list of merchants in it. Each merchant segment yields useful information about the type of merchants, their average purchase and transaction rates, and other statistical information. (Merchant "segments" and merchant "clusters" are used interchangeably herein.)

Brief Summary Text (18):

Preferably, each consumer is also given a profile that includes various demographic data, and summary data on spending habits. In addition, each consumer is preferably given a consumer vector. From the spending data, the merchants that the consumer has most frequently or recently purchased is determined. The consumer vector is then the summation of these merchant vectors. As new purchases are made, the consumer vector is updated, preferably decaying the influence of older purchases. In essence, like the expression "you are what you eat," the present invention reveals "you are whom you shop at," since the vectors of the merchants are used to construct the vectors of the consumers.

Brief Summary Text (20):

Given the merchant segments, the present invention then creates a predictive model of future spending in each merchant segment, based on transaction statistics of historical spending in the merchant segment by those consumers who have purchased from merchants in the segments, in other segments, and data on overall purchases. In one embodiment, each predictive model predicts spending in a merchant cluster in a predicted time interval, such as 3 months, based on historical spending in the cluster in a prior time interval, such as the previous 6 months. During model training, the historical transactions in the merchant cluster for consumers who spent in the cluster, is summarized in each consumer's profile in summary statistics, and input into the predictive model along with actual spending in a predicted time interval. Validation of the predicted spending with actual spending is used to confirm model performance. The predictive models may be a neural networks, or other multivariate statistical model.

Brief Summary Text (22):

To predict financial behavior, the consumer profile of a consumer, using preferably the same type of summary statistics for a recent, past time period, is input into the predictive models for the different merchant clusters. The result is a prediction of the amount of money that the consumer is likely to spend in each merchant cluster in a future time interval, for which no actual spending data may yet be available.

Brief Summary Text (23):

For each consumer, a membership function may be defined which describes how strongly the consumer is associated with each merchant segment. (Preferably, the membership function outputs a membership value for each merchant segment.) The membership function may be the predicted future spending in each merchant segment, or it may be a function of the consumer vector for the consumer and a merchant segment vector (e.g. centroid of each merchant segment). The membership function can be weighted by the amount spent by the consumer in each merchant segment, or other factors. Given the membership function, the merchant clusters for which the consumer has the highest membership values are of particular interest: they are the clusters in which the consumer will spend the most money in the future, or whose spending habits are most similar to the merchants in the cluster. This allows very specific and accurate targeting of promotions, advertising and the like to these consumers. A financial institution using the predicted spending information can direct promotional offers to consumers who are predicted to spend heavily in a merchant segment, with the promotional offers associated with merchants in the merchant segment.

Brief Summary Text (24):

Also, given the membership values, changes in the membership values can be readily determined over time, to identify transitions by the consumer between merchants segments of interest. For example, each month (e.g. after a new credit card billing period or bank statement), the membership function is determined for a consumer, resulting in a new membership value for each merchant cluster. The new membership values can be compared with the previous month's membership values to indicate the largest positive and negative increases, revealing the consumer's changing purchasing habits. Positive changes reflect purchasing interests in new

merchant clusters; negative changes reflect the consumer's lack of interest in a merchant cluster in the past month. Segment transitions such as these further enable a financial institution to target consumers with promotions for merchants in the segments in which the consumers show significant increases in membership values.

Brief Summary Text (25):

In another aspect, the present invention provides an improved methodology for learning the relationships between merchants in transaction data, and defining vectors which represent the merchants. More particularly, this aspect of the invention accurately identifies and captures the patterns of spending behavior which result in the co-occurrence of transactions at different merchants. The methodology is generally as follows:

Brief Summary Text (26):

First, the number of times that each pair of merchants co-occur with one another in the transaction data is determined. The underlying intuition here is that merchants whom the consumers' behaviors indicates as being related will occur together often, whereas unrelated merchants do not occur together often. For example, a new mother will likely shop at children's clothes stores, toy stores, and other similar merchants, whereas a single young male will likely not shop at these types of merchants. The identification of merchants is by counting occurrences of merchants' names in the transaction data. The merchants' names may be normalized to reduce variations and equate different versions of a merchant's name to a single common name.

Brief Summary Text (32):

The present invention may be embodied in various forms. As a computer program product, the present invention includes a data preprocessing module that takes consumer spending data and processes it into organized files of account related and time organized purchases. Processing of merchant names in the spending data is provided to normalize variant names of individual merchants. A data post processing module generates consumer profiles of summary statistics in selected time intervals, for use in training the predictive model. A predictive model generation system creates merchant vectors, and clusters them into merchant clusters, and trains the predictive model of each merchant segment using the consumer profiles and transaction data. Merchant vectors, and consumer profiles are stored in databases. A profiling engine applies consumer profiles and consumer transaction data to the predictive models to provide predicted spending in each merchant segment, and to compute membership functions of the consumers for the merchant segment. A reporting engine outputs reports in various formats regarding the predicted spending and membership information. A segment transition detection engine computes changes in each consumer's membership values to identify significant transitions of the consumer between merchant clusters. The present invention may also be embodied as a system, with the above program product element cooperating with computer hardware components, and as a computer implemented method.

Drawing Description Text (3):

FIG. 2 is a sample list of merchant segments.

Drawing Description Text (6):

FIG. 4b is an illustration of the system architecture of the present invention during development and training of merchant vectors, and merchant segment predictive models.

Drawing Description Text (11):

FIG. 9 is an illustration of the application of multiple consumer account data to the multiple segment predictive models.

Detailed Description Text (1):

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS A. Overview of Consumer and Merchant Vector Representation and the Co-occurrence of Merchant Purchases B. System Overview C. Functional Overview D. Data Preprocessing Module E. Predictive Model Generation System 1. Merchant Vector Generation 2. Training of Merchant Vectors: The UDL Algorithm a) Co-occurrence Counting i) Forward co-occurrence counting ii) Backward co-occurrence counting iii) Bi-directional co-occurrence counting b) Estimating Expected Co-occurrence Counts c) Desired Dot-Products between Merchant Vectors d) Merchant Vector Training 3. Clustering Module F. Data Postprocessing Module G. Predictive Model Generation H. Profiling Engine 1. Membership Function: Predicted Spending In Each Segment 2. Segment Membership Based on Consumer Vectors 3. Updating of Consumer

Profiles I. Reporting Engine 1. Basic Reporting Functionality 2. General Segment Report a) General Segment Information b) Segment Members Information c) Lift Chart d) Population Statistics Tables i) Segment Statistics ii) Row Descriptions J. Targeting Engine K. Segment Transition Detection

Detailed Description Text (3):

One feature of the present invention that enables prediction of consumer spending levels at specific merchants is the ability to represent both consumer and merchants in a same modeling representation. A conventional example is attempting to classify both consumers and merchants with demographic labels (e.g. "baby boomers", or "empty-nesters"). This conventional approach is simply arbitrary, and does not provide any mechanisms for directly quantifying how similar a consumer is to various merchants. The present invention, however, does provide such a quantifiable analysis, based on high-dimensional vector representations of both consumers and merchants, and the co-occurrence of merchants in the spending data of individual consumers.

Detailed Description Text (8):

Thus, in FIG. 1b, following processing of the consumer transaction data, the merchant vectors for merchants A, C, and E have been updated, based on actual spending data, such as C1's transactions, to point generally in the same direction, as have the merchant vectors for merchants B and D, based on C2's transactions. Clustering techniques are used then to identify clusters or segments of merchants based on their merchant vectors 402. In the example of FIG. 1b, a merchant segment is defined to include merchants A, C, and E, such as "upscale-technology_savvy." Note that as defined above, the SIC codes of these merchants are entirely unrelated, and so SIC code analysis would not reveal this group of merchants. Further, a different segment with merchants B and D is identified, even though the merchants share the same SIC codes with the merchants in the first segment, as shown in the transaction data 104.

Detailed Description Text (9):

Each merchant segment is associated with a merchant segment vector 105, preferably the centroid of the merchant cluster. Based on the types of merchants in the merchant segment, and the consumers who have purchased in the segment, a segment name can be defined, and may express the industry, sub-industry, geography, and/or consumer demographics.

Detailed Description Text (10):

The merchant segments provide very useful information about the consumers. In FIG. 1b there is shown the consumer vectors 106 for consumers C1 and C2. Each consumer's vector is a summary vector of the merchants at which the consumer shops. This summary is preferably the vector sum of merchant vectors at which the consumer has shopped at in defined recent time interval. The vector sum can be weighted by the recency of the purchases, their dollar amount, or other factors.

Detailed Description Text (11):

Being in the same vector space as the merchant vectors, the consumer vectors 106 reveal the consumer's interests in terms of their actual spending behavior. This information is by far a better base upon which to predict consumer spending at merchants than superficial demographic labels or categories. Thus, consumer C1's vector is very strongly aligned with the merchant vectors of merchants A, C, and E, indicating C1 is likely to be interested in the products and services of these merchants. C1's vector can be aligned with these merchants, even if C1 never purchased at any of them before. Thus, merchants A, C, and E have a clear means for identifying consumers who may be interested in purchasing from them.

Detailed Description Text (12):

Which consumers are associated with which merchant segments can also be determined by a membership function. This function can be based entirely on the merchant segment vectors and the consumer vectors (e.g. dot product), or on other quantifiable data, such as amount spent by a consumer in each merchant segment, or a predicted amount to be spent.

Detailed Description Text (13):

Given the consumers who are members of a segment, useful statistics can be generated for the segment, such as average amount spent, spending rate, ratios of how much these consumers spend in the segment compared with the population average, and so forth. This information enables merchants to finely target and promote their products to the appropriate consumers.

Detailed Description Text (14):

FIG. 2 illustrates portions of a sample index of merchant segments, as may be produced by the present invention. Segments are named by assigning each segment a unique segment number 200 between 1 and M the total number of segments. In addition, each segment has a description field 210 which describes the merchant segment. A preferred description field is of the form:

Detailed Description Text (15):

Major categories 202 describe how the customers in a merchant segment typically use their accounts. Uses include retail purchases, direct marketing purchases, and where this type cannot be determined, then other major categories, such as travel uses, educational uses, services, and the like. Minor categories 204 describe both a subtype of the major category (e.g. subscriptions being a subtype of direct marketing) or the products or services purchased in the transactions (e.g. housewares, sporting goods, furniture) commonly purchased in the segment.

Demographics information 206 uses account data from the consumers who frequent this segment to describe the most frequent or average demographic features, such as age range or gender, of the consumers. Geographic information 208 uses the account data to describe the most common geographic location of transactions in the segment. In each portion of the segment description 210 one or more descriptors may be used (i.e. multiple major, minor, demographic, or geographic descriptors). This naming convention is much more powerful and fine-grained than conventional SIC classifications, and provides insights into not just the industries of different merchants (as in SIC) but more importantly, into the geographic, approximate age or gender, and lifestyle choices of consumers in each segment.

Detailed Description Text (16):

The various types of segment reports are further described in section I. Reporting Engine, below.

Detailed Description Text (18):

Turning now to FIG. 4a there is shown an illustration of a system architecture of one embodiment of the present invention during operation in a mode for predicting consumer spending. System 400 includes begins with a data preprocessing module 402, a data postprocessing module 410, a profiling engine 412, and a reporting engine 426. Optional elements include a segment transition detection engine 420 and a targeting engine 422. System 400 operates on different types of data as inputs, including consumer summary file 404 and consumer transaction file 406, generates interim models and data, including the consumer profiles in profile database 414, merchant vectors 416, merchant segment predictive models 418, and produces various useful outputs including various segment reports 428-432.

Detailed Description Text (24):

The merchant vectors are then clustered 304 into merchant segments. The merchant segments generally describe groups of merchants which are naturally (in the data) shopped at "together" based on the transactions of the many consumers. Each merchant segment has a segment vector computed for it, which is a summary (e.g. centroid) of the merchant vectors in the merchant segment. Merchant segments provide very rich information about the merchants that are members of the segments, including statistics on rates and volumes of transactions, purchases, and the like.

Detailed Description Text (25):

With the merchant segments now defined, a predictive model of spending behavior is created 306 for each merchant segment. The predictive model for each segment is derived from observations of consumer transactions in two time periods: an input time window and a subsequent prediction time window. Data from transactions in the input time window for each consumer (including both segment specific and cross-segment) is used to extract independent variables, and actual spending in the prediction window provides the dependent variable. The independent variables typically describe the rate, frequency, and monetary amounts of spending in all segments and in the segment being modeled. A consumer vector derived from the consumer's transactions may also be used. Validation and analysis of the segment predictive models may done to confirm the performance of the models.

Detailed Description Text (26):

In the production phase, the system is used to predict spending, either in future time periods for which there is no actual data as of yet, or in a recent past time period for which data is available and which is used for retrospective analysis. Generally, each account (or consumer)

has a profile summarizing the transactional behavior of the account holder. This information is created, or updated 308 with recent transaction data if present, to generate the appropriate variables for input into the predictive models for the segments. (Generation of the independent variables for model generation may also involve updating 308 of account profiles.)

Detailed Description Text (27):

Each account further includes a consumer vector which is derived, e.g. as a summary vector, from the merchant vectors of the merchant at which the consumer has purchased in a defined time period, say the last three months. Each merchant vector's contribution to the consumer vector can be weighted by the consumer's transactions at the merchants, such as by transaction amounts, rates, or recency. The consumer vectors, in conjunction with the merchant segment vectors provide an initial level of predictive power. Each consumer can now be associated with the merchant segment having a merchant segment vector closest to the consumer vector for the consumer.

Detailed Description Text (28):

Using the updated account profiles, this data is input into the set of predictive models to generate 310 for each consumer, an amount of predicted spending in each merchant segment in a desired prediction time period. For example, the predictive models may be trained on a six month input window to predict spending in a subsequent three month prediction window. The predicted period may be an actual future period or a current (e.g. recently ended) period for which actual spending is available.

Detailed Description Text (29):

The predicted spending levels and consumer profiles allow for various levels and types of account and segment analysis 312. First, each account may be analyzed to determine which segment (or segments) the account is a member of, based on various membership functions. A preferred membership function is the predicted spending value, so that each consumer is a member of the segment for which they have the highest predicted spending. Other measures of association between accounts and segments may be based on percentile rankings of each consumer's predicted spending across the various merchant segments. With any of these (or similar) methods of determining which consumers are associated with which segments, an analysis of the rates and volumes of different types of transactions by consumers in each segment can be generated. Further, targeting of accounts in one or more segments may be used to selectively identify populations of consumers with predicted high dollar amount or transaction rates. Account analysis also identifies consumers who have transitioned between segments as indicated by increased or decreased membership values.

Detailed Description Text (30):

Using targeting criteria, promotions directed 314 to specific consumers in specific segments and the merchants in those segments can be realized. For example, given a merchant segment, the consumers with the highest levels (or rankings) of predicted spending in the segment may be identified, or the consumers having consumer vectors closest to the segment vector may be selected. Or, the consumers who have highest levels of increased membership in a segment may be selected. The merchants which make up the segment are known from the segment clustering 304. One or more promotional offers specific to merchants in the segment can be created, such as discounts, incentives and the like. The merchant-specific promotional offers are then directed to the selected consumers. Since these account holders have been identified as having the greatest likelihood of spending in the segment, the promotional offers beneficially coincide with their predicted spending behavior. This desirably results in an increase success rate at which the promotional offers are redeemed.

Detailed Description Text (33):

The data preprocessing module 402 (DPM) does initial processing of consumer data received from a source of consumer accounts and transactions, such as a credit card issuer, in preparation for creating the merchant vectors, consumer vectors, and merchant segment predictive models. DPM 402 is used in both production and training modes. (In this disclosure, the terms "consumer," "customer," and "account holder" are used interchangeably).

Detailed Description Text (35):

Customer summary file 404: The customer summary file 404 contains one record for each customer that is profiled by the system, and includes account information of the customer's account, and optionally includes demographic information about the customer. The consumer summary file 404

is typically one that a financial institution, such as a bank, credit card issuer, department store, and the like maintains on each consumer. The customer or the financial institution may supply the additional demographic fields which are deemed to be of informational or of predictive value. Examples of demographic fields include age, gender and income; other demographic fields may be provided, as desired by the financial institution.

Detailed Description Text (36):

Table 1 describes one set of fields for the customer summary file 404 for a preferred embodiment. Most fields are self-explanatory. The only required field is an account identifier that uniquely identifies each consumer account and transactions. This account identifier may be the same as the consumer's account number; however, it is preferable to have a different identifier used, since a consumer may have multiple account relationships with the financial institution (e.g. multiple credit cards or bank accounts), and all transactions of the consumer should be dealt with together. The account identifier is preferably derived from the account number, such as by a one-way hash or encrypted value, such that each account identifier is uniquely associated with an account number. The pop_id field is optionally used to segment the population of customers into arbitrary distinct populations as specified by the financial institution, for example by payment history, account type, geographic region, etc.

Detailed Description Text (37):

Note the additional, optional demographic fields for containing demographic information about each consumer. In addition to demographic information, various summary statistics of the consumer's account may be included. These include any of the following:

Detailed Description Text (44):

The DPM 402 creates the master file 408 from the consumer summary file 404 and consumer transaction file 406 by the following process: a) Verify minimum data requirements. The DPM 402 determines the number of data files it is handling (since there maybe many physical media sources), and the length of the files to determine the number of accounts and transactions. Preferably, a minimum of 12 months of transactions for a minimum of 2 million accounts are used to provide fully robust models of merchants and segments. However, there is no formal lower bound to the amount of data on which system 400 may operate. b) Data cleaning. The DPM 402 verifies valid data fields, and discards invalid records. Invalid records are records that are missing the any of the required fields for the customer summary file of the transaction file. The DPM 402 also indicates missing values for fields that have corrupt or missing data and are optional. Duplicate transactions are eliminated using account ID, account number, transaction code, transaction amount, date, and merchant description as a key. c) Sort and merge files. The consumer summary file 404 and the consumer transaction file 406 are both sorted by account ID; the consumer transaction file 406 is further sorted by transaction date. Additional sorting of the transaction file, for example on time, type of transaction, merchant zip code, may be applied to further influence the determination of merchant co-occurrence. The sorted files are merged into the master file 408, with one record per account, as described above.

Detailed Description Text (47):

Referring to FIG. 4b, the predictive model generation system 440 takes as its inputs the master file 408 and creates the consumer profiles and consumer vectors, the merchant vectors and merchant segments, and the segment predictive models. This data is used by the profiling engine to generate predictions of future spending by a consumer in each merchant segment using inputs from the data postprocessing module 410.

Detailed Description Text (118):

The second technique, UDL2, overcomes of the small count problem by using log-likelihood ratio estimates to calculate $r_{sub.ij}$. It has been shown that log-likelihood ratios have much better small count behavior than χ^2 , while at the same time retaining the same behavior as χ^2 in the non-small count regions.

Detailed Description Text (136):

Following generation and training of the merchant vectors, the clustering module 520 is used to cluster the resulting merchant vectors and identify the merchant segments. Various different clustering algorithms may be used, including k-means clustering (MacQueen). The output of the clustering is a set of merchant segment vectors, each being the centroid of a merchant segment, and a list of merchant vectors (thus merchants) included in the merchant segment.

Detailed Description Text (137):

There are two different clustering approaches that may be usefully employed to generate the merchant segments. First, clustering may be done on the merchant vectors themselves. This approach looks for merchants having merchant vectors which are substantially aligned in the vector space, and clusters these merchants into segments and computes a cluster vector for each segment. Thus, merchants for whom transactions frequently co-occur and have high dot products between their merchant vectors will tend to form merchant segments. Note that it is not necessary for all merchants in a cluster to all co-occur in many consumers' transactions. Instead, co-occurrence is associative: if merchants A and B co-occur frequently, and merchants B and C co-occur frequently, A and C are likely to be in the same merchant segment.

Detailed Description Text (142):

However computed, the consumer vectors can then be clustered, so that similar consumers, based on their purchasing behavior, form a merchant segment. This defines a merchant segment vector. The merchant vectors which are closest to a particular merchant segment vector are deemed to be included in the merchant segment.

Detailed Description Text (143):

With the merchant segments and their segment vectors, the predictive models for each segment may be developed. Before discussing the creation of the predictive models, a description of the training data used in this process is described.

Detailed Description Text (145):

Following identification of merchant segments, a predictive model of consumer spending in each segment is generated from past transactions of consumers in the merchant segment. Using the past transactions of consumer in the merchant segment provides a robust base on which to predict future spending, and since the merchant segments were identified on the basis of the actual spending patterns of the consumers, the arbitrariness of conventional demographic based predictions are minimized. Additional non-segment specific transactions of the consumer may also be used to provide a base of transaction behavior.

Detailed Description Text (146):

To create the segment models, the consumer transaction data is organized into groups of observations. Each observation is associated with a selected end-date. The end-date divides the observation into a prediction window and an input window. The input window includes a set of transactions in a defined past time interval prior to the selected end-date (e.g. 6 months prior). The prediction window includes a set of transactions in a defined time interval after the selected end-date (e.g. the next 3 months). The prediction window transactions are the source of the dependent variables for the prediction, and the input window transactions are the source of the independent variables for the prediction.

Detailed Description Text (148):

The first type of observations are training observations which are used to train the predictive models that predicts future spending within particular merchant segments. If N is the length (in months) of the window over which observation inputs are computed then there are $2N-1$ training observations for each segment.

Detailed Description Text (150):

The second type of observations are blind observations. Blind observations are observations where the prediction window does not overlap any of the time frames for the prediction windows in the training observations. Blind observations are used to evaluate segment model performance. In FIG. 8, the blind observations 804 include those from September to February, as illustrated.

Detailed Description Text (152):

FIG. 8 also illustrates that at some point during the prediction window, the financial institution sends out promotions to selected consumers based on their predicted spending in the various merchant segments.

Detailed Description Text (153):

Referring to FIG. 4b again, the DPPM takes the master files 408, and a given selected end-date, and constructs for each consumer, and then for each segment, a set of training observations and blind observations from the consumer's transactions, including transactions in the segment, and

any other transactions. Thus, if there are 300 segments, for each consumer there will be 300 sets of observations. If the DPPM is being used during production for prediction purposes, then the set of observations is a set of action observations.

Detailed Description Text (155):

Prediction window: The dependent variables are generally any measure of amount or rate of spending by the consumer in the segment in the prediction window. A simple measure is the total dollar amount that was spent in the segment by the consumer in the transactions in the prediction window. Another measure may be average amount spent at merchants (e.g. total amount divided by number of transactions).

Detailed Description Text (156):

Input window: The independent variables are various measures of spending in the input window leading up to the end date (though some may be outside of it). Generally, the transaction statistics for a consumer can be extracted from various grouping of merchants. These groups may be defined as: 1) merchants in all segments; 2) merchants in the merchant segment being modeled; 3) merchants whose merchant vector is closest the segment vector for the segment being modeled (these merchants may or may not be in the segment); and 4) merchants whose merchant vector is closest to the consumer vector of the consumer.

Detailed Description Text (157):

One preferred set of input variables includes: (1) Recency. The amount of time in months between the current end date and the most recent transaction of the consumer in any segment. Recency may be computed over all available time and is not restricted to the input window. (2) Frequency. The number of transactions by a consumer in the input window preceding the end-date for all segments. (3) Monetary value of purchases. A measure of the amount of dollars spent by a customer in the input window preceding the end-date for all segments. The total or average, or other measures may be used. (4) Recency segment. The amount of time in months between the current end date and the most recent transaction of the consumer in the segment. Recency may be computed over all available time and is not restricted to the input window. (5) Frequency segment. The number of transactions in the segment by a customer in the input window preceding the current end date. (6) Monetary segment. The amount of dollars spent in the segment by a customer in the input window preceding the current end date. (7) Recency nearest profile merchants. The amount of time in months between the current end date and the most recent transaction of the consumer in a collection of merchants that are nearest the consumer vector of the consumer. Recency may be computed over all available time and is not restricted to the input window. (8) Frequency nearest profile merchants. The number of transactions in a collection of merchants that are nearest the consumer vector of the consumer by the consumer in the input window preceding the current end date. (9) Monetary nearest frequency merchants. The amount of dollars spent in a collection of merchants that are nearest the consumer vector of the consumer in the input window preceding the current end date. (10) Recency nearest segment merchants. The amount of time in months between the current end date and the most recent transaction of the consumer in a collection of merchants that are nearest the segment vector. Recency may be computed over all available time and is not restricted to the input window. (11) Frequency nearest segment merchants. The number of transactions in a collection of merchants that are nearest the segment vector by the consumer in the input window preceding the current end date. (12) Monetary nearest segment merchants. The amount of dollars spent in a collection of merchants that are nearest the segment vector by the consumer in the input window preceding the current end date. (13) Segment probability score. The probability that a consumer will spend in the segment in the prediction window given all merchant transactions for the consumer in the input window preceding the end date. A preferred algorithm estimates combined probability using a recursive Bayesian method. (14) Seasonality variables. It is assumed that the fundamental period of the cyclic component is known. In the case of seasonality, it can be assumed that the cycle of twelve months. Two variables are added to the model related to seasonality. The first variable codes the sine of the date and the second variable codes the cosine of the date. The calculation for these variables are:

Detailed Description Text (158):

In addition to these transaction statistics, variables may be defined for the frequency of purchase and monetary value for all cases of segment merchants, nearest profile merchants, nearest segment merchants for the same forward prediction window in the previous year(s).

Detailed Description Text (160):

The training observations for each segment are input into the segment predictive model generation module 530 to generate a predictive model for the segment. FIG. 9 illustrates the overall logic of the predictive model generation process. The master files 408 are organized by accounts, based on account identifiers, here illustratively, accounts 1 through N. There are M segments, indicated by segments 1 through M. The DPPM generates for each combination of account and merchant segment, a set of input and blind observations. The respective observations for each merchant segment M from the many accounts 1 . . . N are input into the respective segment predictive model M during training. Once trained, each segment predictive model is tested with the corresponding blind observations. Testing may be done by comparing for each segment a lift chart generated by the training observations with the lift chart generated from blind observations. Lift charts are further explained below.

Detailed Description Text (164):

The profiling engine 412 provides analytical data in the form of an account profile about each customer whose data is processed by the system 400. The profiling engine is also responsible for updating consumer profiles over time as new transaction data for consumers is received. The account profiles are objects that can be stored in a database 414 and are used as input to the computational components of system 400 in order to predict future spending by the customer in the merchant segments. The profile database 414 is preferably ODBC compliant, thereby allowing the accounts provider (e.g. financial institution) to import the data to perform SQL queries on the customer profiles.

Detailed Description Text (165):

The account profile preferably includes a consumer vector, a membership vector describing a membership value for the consumer for each merchant segment, such as the consumer's predicted spending in each segment in a predetermined future time interval, and the recency, frequency, and monetary variables as previously described for predictive model training.

Detailed Description Text (166):

The profiling engine 412 creates the account profiles as follows.

Detailed Description Text (167):

1. Membership Function: Predicted Spending in Each Segment

Detailed Description Text (168):

The profile of each account holder includes a membership value with respect to each segment. The membership value is computed by a membership function. The purpose of the membership function is to identify the segments with which the consumer is mostly closely associated, that is, which best represent the group or groups of merchants at which the consumer has shopped, and is likely to shop at in the future.

Detailed Description Text (169):

In a preferred embodiment, the membership function computes the membership value for each segment as the predicted dollar amount that the account holder will purchase in the segment given previous purchase history. The dollar amount is projected for a predicted time interval (e.g. 3 months forward) based on a predetermined past time interval (e.g. 6 months of historical transactions). These two time intervals correspond to the time intervals of the input window and prediction windows used during training of the merchant segment predictive models. Thus, if there are 300 merchant segments, then a membership value set is a list of 300 predicted dollar amounts, corresponding to the respective merchant segments. Sorting the list by the membership value identifies the merchant segments at which the consumer is predicted to spend the greatest amounts of money in the future time interval, given their spending historically.

Detailed Description Text (170):

To obtain the predicted spending, certain data about each account is input in each of the segment predictive models. The input variables are constructed for the profile consistent with the membership function of the profile. Preferably, the input variables are the same as those used during model training, as set forth above. An additional input variable for the membership function may include the dot product between the consumer vector and the segment vector for the segment (if the models are so trained). The output of the segment models is a predicted dollar amount that the consumer will spend in each segment in the prediction time interval.

Detailed Description Text (171):2. Segment Membership Based on Consumer VectorsDetailed Description Text (172):

A second alternate, membership aspect of the account profiles is membership based upon the consumer vector for each account profile. The consumer vector is a summary vector of the merchants that the account has shopped at, as explained above with respect to the discussion of clustering. In this aspect, the dot product of the consumer vector and segment vector for the segment defines a membership value. In this embodiment, the membership value list is a set of 300 dot products, and the consumer is member of the merchant segment(s) having the highest dot product(s).

Detailed Description Text (173):

With either one of these membership functions, the population of accounts that are members of each segment (based on the accounts having the highest membership values for each segment) can be determined. From this population, various summary statistics about the accounts can be generated such as cash advances, purchases, debits, and the like. This information is further described below.

Detailed Description Text (174):3. Updating of Consumer ProfilesDetailed Description Text (180):

The reporting engine 426 provides various types of segment and account specific reports. The reports are generated by querying the profiling engine 412 and the account database for the segments and associated accounts, and tabulating various statistics on the segments and accounts.

Detailed Description Text (184):2. General Segment ReportDetailed Description Text (185):

For each merchant segment a very detailed and powerful analysis of the segment can be created in a segment report. This information includes:

Detailed Description Text (186):a) General Segment InformationDetailed Description Text (187):

Merchant Cohesion: A measure of how closely clustered are the merchant vectors in this segment. This is the average of the dot products of the merchant vectors with the centroid vector of this segment. Higher numbers indicate tighter clustering.

Detailed Description Text (188):

Number of Transactions: The number of purchase transactions at merchants in this segment, relative to the total number of purchase transactions in all segments, providing a measure of how significant the segment is in transaction volume.

Detailed Description Text (189):

Dollars Spent: The total dollar amount spent at merchants in this segment, relative to the total dollar amount spent in all segments, providing a measure of dollar volume for the segment.

Detailed Description Text (190):

Most Closely Related Segments: A list of other segments that are closest to the current segment. This list may be ranked by the dot products of the segment vectors, or by a measure of the conditional probability of purchase in the other segment given a purchase in the current segment.

Detailed Description Text (191):

The conditional probability measure M is as follows: $P(A.\text{vertline}.B)$ is probability of purchase in segment A segment in next time interval (e.g. 3 months) given purchases in segment B in the previous time interval (e.g. 6 months). $P(A.\text{vertline}.B)/P(A)=M$. If M is >1, then a purchase in

segment B is positively influencing the probability of purchase in segment A, and if $M < 1$ then a purchase in segment B negatively influences a purchase in segment A. This is because if there is no information about the probability of purchases in segment B, then $P(A.\text{vertline.}B) = P(A)$, so $M = 1$. The values for $P(A.\text{vertline.}B)$ are determined from the co-occurrences of purchases at merchants in the two segments, and $P(A)$ is determined and from the relative frequency of purchases in segment A compared to all segments.

Detailed Description Text (192):

A farthest segments list may also be provided (e.g. with the lowest conditional probability measures).

Detailed Description Text (193):

b) Segment Members Information

Detailed Description Text (194):

Detailed information is provided about each merchant which is a member of a segment. This information comprises: Merchant Name and SIC code; Dollar Bandwidth: The fraction of all the money spent in this segment that is spent at this merchant (percent); Number of transactions: The number of purchase transactions at this merchant; Average Transaction Amount: The average value of a purchase transaction at this merchant; Merchant Score: The dot product of this merchant's vector with the centroid vector of the merchant segment. (A value of 1.0 indicates that the merchant vector is at the centroid); SIC Description: The SIC code and its description;

Detailed Description Text (198):

Table 10 illustrates a sample lift chart for merchant segment:

Detailed Description Text (201):

For each merchant segment then, the consumer accounts are ranked by their predicted spending for the segment in the prediction window period. Once the accounts are ranked, they are divided into N (e.g. 20) equal sized bins so that bin 1 has the highest spending accounts, and bin N has the lowest ranking accounts. This identifies the account holders that the predictive model for the segment indicated should be expected to spend the most in this segment.

Detailed Description Text (202):

Then, for each bin, the average actual spending per account in this segment in the past time period, and the average predicted spending is computed. The average actual spending over all bins is also computed. This average actual spending for all accounts is the baseline spending value (in dollars), as illustrated in the last line of Table 10. This number describes the average that all account holders spent in the segment in the prediction window period.

Detailed Description Text (203):

The lift for a bin is the average actual spending by accounts in the bin divided by the baseline spending value. If the predictive model for the segment is accurate, then those accounts in the highest ranked bins should have a lift greater than 1, and the lift should generally be increasing, with bin 1 having the highest lift. Where this is the case, as for example, in Table 10, in bin 1, this shows that those accounts in bin 1 in fact spent several times the baseline, thereby confirming the prediction that these accounts would in fact spend more than others in this segment.

Detailed Description Text (205):

The lift information allows the financial institution to very selectively target a specific group of accounts (e.g. the accounts in bin 1) with promotional offers related to the merchants in the segment. This level of detailed, predictive analysis of very discrete groups of specific accounts relative to merchant segments is not believed to be currently available by conventional methods.

Detailed Description Text (207):

The reporting engine 426 further provides two types of analyses of the financial behavior of a population of accounts that are associated with a segment based on various selection criteria. The Segment Predominant Scores Account Statistics table and the Segment Top 5% Scores Account Statistics table present averaged account statistics for two different types of populations of customers who shop, or are likely to shop, in a given segment. The two populations are

determined as follows.

Detailed Description Text (208):

Segment Predominant Scores Account Statistics Table: All open accounts with at least one purchase transaction are scored (predicted spending) for all of the segments. Within each segment, the accounts are ranked by score, and assigned a percentile ranking. The result is that for each account there is a percentile ranking value for each of the merchant segments.

Detailed Description Text (209):

The population of interest for a given segment is defined as those accounts which have their highest percentile ranking in this segment. For example, if an account has its highest percentile ranking in segment #108, that account will be included in the population for the statistics table for segment #108, but not in any other segment. This approach assigns each account holder to one and only one segment.

Detailed Description Text (210):

Segment Top 5% Scores Account Statistics. For the Segment Top 5% Scores Account Statistics table, the population is defined as the accounts with percentile ranking of 95% or greater in a current segment. These are the 5% of the population that is predicted to spend the most in the segment in the predicted future time interval following the input data time window. These accounts may appear in this population in more than one segment, so that high spenders will show up in many segments; concomitantly, those who spend very little may not be assigned to any segment.

Detailed Description Text (213):

i) Segment Statistics

Detailed Description Text (214):

The tables present the following statistics for each of several categories, one category per row. The statistics are: Mean Value: the average over the population being scored; Std Deviation: the standard deviation over the population being scored; Population Mean: the average, over all the segments, of the Mean Value (this column is thus the same for all segments, and are included for ease of comparison); and Relative Score: the Mean Value, as a fraction of the Population Mean (in percent).

Detailed Description Text (216):

Each table contains rows for spending and rate in Cash Advances, Purchases, Debits, and Total Spending. The rows for spending (Cash Advances, Purchases, and Debits) show statistics on dollars per month for all accounts in the population over the time period of available data. The rate rows (Cash Advance Rate, Debit Rate, and Purchase Rate) show statistics on the number of transactions per month for all accounts in the population over the time period of available data. Debits consist of Cash Advances and Purchases. The Dollars in Segment shows the fraction of total spending that is spent in this segment. This informs the financial institution of how significant overall this segment is. The Rate in Segment shows the fraction of total purchase transactions that occur in this segment.

Detailed Description Text (217):

The differences between these two populations are subtle but important, and are illustrated by the above tables. The segment predominant population identifies those individuals as members of a segment who, relative to their own spending, are predicted to spend the most in the segment. For example, assume a consumer whose predicted spending in a segment is \$20.00, which gives the consumer a percentile ranking of 75.sup.th percentile. If the consumer's percentile ranking in every other segment is below the 75.sup.th percentile, then the consumer is selected in this population for this segment. Thus, this may be considered an intra-account membership function.

Detailed Description Text (218):

The Top 5% scores population instead includes those accounts holders predicted to spend the most in the segment, relative to all other account holders. Thus, the account holder who was predicted to spend only \$20.00 in the merchant segment will not be member of this population since he is well below the 95.sup.th percentile, which may be predicted to spend, for example \$100.00.

Detailed Description Text (219):

In the example tables these differences are pronounced. In Table 11, the average purchases of the segment predominant population is only \$166.86. In Table 12, the average purchase by top 5% population is more than twice that, at \$391.54. This information allows the financial institution to accurately identify accounts which are most likely to spend in a given segment, and target these accounts with promotional offers for merchants in the segment.

Detailed Description Text (220):

The above tables may also be constructed based on other functions to identify accounts associated with segments, including dot products between consumer vectors and segment vectors.

Detailed Description Text (222):

The targeting engine 422 allows the financial institution to specify targeted populations for each (or any) merchant segment, to enable selection of the targeted population for receiving predetermined promotional offers.

Detailed Description Text (223):

A financial institution can specify a targeted population for a segment by specifying a population count for the segment, for example, the top 1000 accounts holders, or the top 10% account holders in a segment. The selection is made by any of the membership functions, including dot product, or predicted spending. Other targeting specifications may be used in conjunction with these criteria, such as a minimum spending amount in the segment, such as \$100. The parameters for selecting the targeting population are defined in a target specification document 424 which is an input to the targeting engine 422. One or more promotions can be specifically associated with certain merchants in a segment, such as the merchants with the highest correlation with the segment vector, highest average transaction amount, or other selective criteria. In addition, the amounts offered in the promotions can be specific to each consumer selected, and based on their predicted or historical spending in the segment. The amounts may also be dependent on the specific merchant for whom a promotion is offered, as a function of the merchant's contributions to purchases in the segment, such as based upon their dollar bandwidth, average transaction amount, or the like.

Detailed Description Text (224):

The selected accounts can be used to generate a targeted segmentation report 430 by providing the account identifiers for the selected accounts to the reporting engine 426, which constructs the appropriate targeting report on the segment. This report has the same format as the general segment report but is compiled for the selected population.

Detailed Description Text (226):

Table 13 shows a specification of a total of at least 228,000 customer accounts distributed over four segments and two promotional offers (ID 1 and ID 2). For each segment or promotional offer, there are different selection and filtering criteria. For promotion #1 the top 75,000 consumers in segment #122 based on predicted spending, and who have an average transaction in the segment greater than \$50, are selected. For this promotion in segment #413, the top 10% of accounts based on the dot product between the consumer vector and segment vector are selected, so long as they have a minimum spending in the segment of \$100. Finally, for promotion #2, 87,000 consumers are selected across two segments. Within each offer (e.g. offer ID 1) the segment models may be merged to produce a single lift chart which reflects the offer as a composition of the segments. The targeting engine 422 then provides the following additional functionality: 1. Select fields from the account profile of the selected accounts that will be inserted to the mail file 434. For example, the name, address, and other information about the account may be extracted. 2. The mail file 434 is then exported to a useful word processing or bulk mailing system. 3. Instruct the reporting engine 426 to generate reports that have summary and cumulative frequencies for select account fields, such as including purchase, debit, cash advance, or any other account data. 4. Instruct the reporting engine 426 to generate lift charts for the targeting population in the segment, and for overlapped (combined) segments.

Detailed Description Text (227):

K. Segment Transition Detection

Detailed Description Text (228):

As is now apparent, the system of the present invention provides detailed insight into which merchant segments a consumer is associated with based on various measures of membership, such

as dot product, predicted spending, and the like. Further, since the consumers continue to spend over time, the consumer accounts and the consumers' associations with segments is expected to change over time as their individual spending habits change.

Detailed Description Text (229):

The present invention allows for detection of the changes in consumer spending via the segment transition detection engine 420. In a given data period (e.g. next monthly cycle or multiple month collection of data) a set of membership values for each consumer is defined as variously described above, with respect to each segment. Again, this may be predicted spending by the consumer in each segment, dot product between the consumer vector and each segment vectors, or other membership functions.

Detailed Description Text (230):

In a subsequent time interval, using additional spending and/or predicted data, the membership values are recomputed. Each consumer will have the top P and the bottom Q increases in and decreases in segment membership. That is, there will be two changes of interest: the P (e.g. 5) segments with the greatest increase in membership values for the consumer; the Q segments with the greatest decrease in segment membership.

Detailed Description Text (231):

An increase in the membership value for a segment indicates that the consumer is now (or predicted to) spend more money in a particular segment. Decreases show a decline in the consumer's interest in the segment. Either of these movements may reflect a change in the consumer's lifestyle, income, or other demographic factors.

Detailed Description Text (232):

Significant increases in merchant segments which previously had low membership values are particularly useful to target promotional offers to the account holders who are moving into the segment. This is because the significant increase in membership indicates that the consumer is most likely to be currently receptive to the promotional offers for merchants in the segment, since they are predicted to be purchasing more heavily in the segment.

Detailed Description Text (233):

Thus, the segment transition detection engine 420 calculates the changes in each consumer's membership values between two selected time periods, typically using data in a most recent prediction window (either ending or beginning with a current statement date) relative to memberships in prior time intervals. The financial institution can define a threshold change value for selecting accounts with changes in membership more significant than the threshold. The selected accounts may then be provided to the reporting engine 426 for generation of various reports, including a segment transition report 432 which is like the general segment report except that it applies to accounts that are considered to have transitioned to or from a segment. This further enables the financial institution to selectively target these customers with promotional offers for merchants in the segments in which the consumer had the most significant positive increases in membership.

Detailed Description Text (234):

In summary then, the present invention provides a variety of powerful analytical methods which predict consumer financial behavior in discretely defined merchant segments, and with respect to predetermined time intervals. The clustering of merchants in merchant segments allows analysis of transactions of consumers in each specific segment, both historically, and in the predicted period to identify consumers of interest. Identified consumers can then be targeted with promotional offers precisely directed at merchants within specific segments.

Detailed Description Paragraph Equation (1):

Major Categories: Minor Categories: Demographics: Geography

Detailed Description Paragraph Equation (10):

$\text{Cos Input} = \text{cos}(2.0 * \text{PI} * (\text{sample month of year}) / 365)$. (15) (Segment Vector-Consumer Vector Closeness: As an optional input, the dot product of the segment vector for the segment and the consumer vector is used as an input variable.

Detailed Description Paragraph Table (1):

TABLE 1 Customer Summary File Description Sample Format Account_id Char[max 24] Pop_id Char

(`1`-'N`) Account number Char[max 16] Credit bureau score Short int as string Internal credit risk Short int as string score Ytd purchases Int as string Ytd_cash_adv Int as string Ytd_int purchases Int as string Ytd_int cash_adv Int as string State code Char[max 2] Zip_code Char[max 5] Demographic 1 Int as string . . . Demographic N Int as string

Detailed Description Paragraph Table (2):

TABLE 2 Example Demographic Fields for Customer Summary File Description Explanation Cardholder zip code Months on books or open date Number of people on the account Equivalent to number of plastics Credit risk score Cycles delinquent Credit line Open to buy Initial month statement balance Balance on the account prior to the first month of transaction data pull Last month statement balance Balance on the account at the end of the transaction data pulled Monthly payment amount For each month of transaction data contributed or the average over last year. Monthly cash advance amount For each month of transaction data contributed or the average over last year. Monthly cash advance count For each month of transaction data contributed or the average over last year. Monthly purchase amount For each month of transaction data contributed or the average over last year. Monthly purchase count For each month of transaction data contributed or the average over last year. Monthly cash advance interest For each month of transaction data contributed or the average over last year. Monthly purchase interest For each month of transaction data contributed or the average over last year. Monthly late charge For each month of transaction data contributed or the average over last year.

Detailed Description Paragraph Table (4):

TABLE 4 Master File 408 Description Sample Format Account id Char[max 24] Pop_id Char (`1`-'N`) Account number Char[max 16] Credit bureau score Short int as string Ytd purchases Int as string Ytd_cash_advances Int as string Ytd_interest on purchases Int as string Ytd_interest on cash advs Int as string State_code Char[max 2] Demographic 1 Int as string . . . Demographic N Int as string <transactions>

Detailed Description Paragraph Table (10):

TABLE 10 A sample segment lift chart Cumulative Cumulative Cumulative Bin segment lift segment lift in S Population 1 5.56 \$109.05 50,000 2 4.82 \$94.42 100,000 3 3.82 \$74.92 150,000 4 3.23 \$63.38 200,000 5 2.77 \$54.22 250,000 6 2.43 \$47.68 300,000 7 2.20 \$43.20 350,000 8 2.04 \$39.98 400,000 9 1.88 \$36.79 450,000 10 1.75 \$34.35 500,000 11 1.63 \$31.94 550,000 12 1.52 \$29.75 600,000 13 1.43 \$28.02 650,000 14 1.35 \$26.54 700,000 15 1.28 \$25.08 750,000 16 1.21 \$23.81 800,000 17 1.16 \$22.65 850,000 18 1.10 \$21.56 900,000 19 1.05 \$20.57 950,000 20 1.00 \$19.60 1,000,000 Base-line -- \$19.60

Detailed Description Paragraph Table (11):

TABLE 11 Segment Predominant Scores Account Statistics: 8291 accounts (0.17 percent) Mean Std Population Relative Category Value Deviation Mean Score Cash Advances \$11.28 \$53.18 \$6.65 169.67 Cash Advance Rate 0.03 0.16 0.02 159.92 Purchases \$166.86 \$318.86 \$192.91 86.50 Purchase Rate 0.74 1.29 1.81 40.62 Debits \$178.14 \$324.57 \$199.55 89.27 Debit Rate 0.77 1.31 1.84 41.99 Dollars in Segment 4.63 14.34 10.63% 43.53 Rate in Segment 3.32 9.64 11.89% 27.95

Detailed Description Paragraph Table (12):

TABLE 11 Segment Predominant Scores Account Statistics: 8291 accounts (0.17 percent) Mean Std Population Relative Category Value Deviation Mean Score Cash Advances \$11.28 \$53.18 \$6.65 169.67 Cash Advance Rate 0.03 0.16 0.02 159.92 Purchases \$166.86 \$318.86 \$192.91 86.50 Purchase Rate 0.74 1.29 1.81 40.62 Debits \$178.14 \$324.57 \$199.55 89.27 Debit Rate 0.77 1.31 1.84 41.99 Dollars in Segment 4.63 14.34 10.63% 43.53 Rate in Segment 3.32 9.64 11.89% 27.95

Detailed Description Paragraph Table (13):

TABLE 13 Target population specification ID associated with Customer promotional Segment target Selection offer ID count Criteria Filter Criteria 1 122 75,000 Predicted Average Spending in Transaction in Segment >\$50 1 143 Top 10% Dot Product Total Spending in Segment >\$100 2 12 and 55 87,000 Predicted None Spending in this Segment 12 and 55

Other Reference Publication (2):

Phillips Business Information. HNC System Sheds New Light on Cardholder Profile. Potomac: Card News, Dec. 7, 1998, v13, n23, p. 4-5.*

CLAIMS:

1. A method of predicting financial behavior of consumers, comprising: generating from transaction data for a plurality of consumers, a date ordered sequence of transactions for each consumer; selecting for each consumer a set of the date ordered transactions to form a group of input transactions for the consumer; and for each consumer, applying the input transactions of the consumer to each of a plurality of merchant segment predictive models, each merchant segment predictive model defining for a group of merchants a prediction function between input transactions in a past time interval and predicted spending in a subsequent time interval, to produce for each consumer a predicted spending amount in each merchant segment.

2. The method of claim 1, further comprising: for each consumer, associating the consumer with the merchant segment for which the consumer had the highest predicted spending relative to other merchant segments.

3. The method of claim 1, further comprising: for each merchant segment, determining a segment vector as a summary vector of merchant vectors of merchants associated with the segment; and for each consumer, associating the consumer with the merchant segment having the greatest dot product between the segment vector of the segment and a consumer vector of the consumer.

4. The method of claim 1, further comprising: for each merchant segment: ranking the consumers by their predicted spending in the merchant segment; determining for each consumer a percentile ranking in the merchant segment; for each consumer: determining the merchant segment in which the consumer's percentile ranking is the highest, to uniquely associate each consumer with one merchant segment; and for each merchant segment, determining summary transaction statistics for the consumers uniquely associated with the merchant segment.

5. The method of claim 1, further comprising: for each merchant segment: ranking the consumers by their predicted spending in the merchant segment; determining for each consumer a percentile ranking in the merchant segment; selecting as a population, the consumers having a percentile ranking in excess of predetermined percentile threshold; and determining summary transaction statistics for selected population of consumers.

14. The method of claim 1, further comprising: determining for each merchant name in the transaction data a merchant vector; clustering the merchant vectors to form a plurality of merchant segments, wherein each merchant vector is associated with one and only one merchant segment; for each merchant segment, determining from the transactions of consumers at the associated merchants of the merchant, statistical measures of consumer transactions in the segment.

15. The method of claim 1, further comprising: selecting a plurality of consumers associated with at least one merchant segment, the selected plurality selected according to their predicted spending in the merchant segment; and providing promotional offers to the selected plurality of consumers.

16. The method of claim 1, further comprising: training each of the merchant segment predictive models to predict spending in a predicted time period based upon transaction statistics of the consumer's transactions in a past time period.

17. The method of claim 16, wherein the transaction statistics comprises variables describing the recency of the consumer's transactions in one or more merchant segments, the frequency of the consumer's transactions in one or more merchant segments, and the amount of the consumer's transactions in one or more merchant segments.

18. A system for predicting consumer financial behavior, comprising: a plurality of merchant segments, each merchant segment having a set of merchants associated therewith; a plurality of merchant segment predictive models, each model associated with one of the merchant segments for predicting spending by an individual consumer in the merchant segment in a predicted time period as a function of transaction statistics of the consumer for transactions in a prior time period; and a data processing module that receives transaction data for a consumer, and constructs the transaction statistics for the prior time period for input into selected ones of the merchant segment predictive models.

19. A system for forming merchant segments, comprising: a data processing module that receives consumer transaction data for a plurality of consumer accounts, and organizes the transaction

data by account, and within account, sequences the transactions by time; a data processing module that determines from the sequenced transaction data an expected frequency of co-occurrence for each merchant, and that constructs for each merchant a merchant vector as a function of unexpected frequency of co-occurrences of the merchant; and a clustering module that clusters the merchant vectors into merchant segment by determining merchant vectors that are closely aligned with each other.

[Previous Doc](#)[Next Doc](#)[Go to Doc#](#)

Refine Search

Search Results -

Term	Documents
TREE	153525
TREES	69260
STRUCTURE	4421164
STRUCTURES	1113747
(6 AND (TREE NEAR STRUCTURE)).PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD.	0
(L6 AND (TREE NEAR STRUCTURE)).PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD.	0

Database:

US Pre-Grant Publication Full-Text Database
 US Patents Full-Text Database
 US OCR Full-Text Database
 EPO Abstracts Database
 JPO Abstracts Database
 Derwent World Patents Index
 IBM Technical Disclosure Bulletins

Search:

Refine Search

Recall Text

Clear

Interrupt

Search History

DATE: Monday, May 23, 2005 [Printable Copy](#) [Create Case](#)

<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>	<u>Set Name</u>
			result set
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>			
<u>L8</u>	16 and (tree near structure)	0	<u>L8</u>
<u>L7</u>	L6 and partition\$	1	<u>L7</u>
<u>L6</u>	L5 and behavior	7	<u>L6</u>
<u>L5</u>	L4 and demographic\$	7	<u>L5</u>
<u>L4</u>	L3 and segment	7	<u>L4</u>
<u>L3</u>	L1 and (profile)	9	<u>L3</u>
<u>L2</u>	L1 and (profile near definition)	1	<u>L2</u>
<u>L1</u>	consumer near cluster\$	33	<u>L1</u>

END OF SEARCH HISTORY

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L6: Entry 6 of 7

File: USPT

Jan 4, 2005

DOCUMENT-IDENTIFIER: US 6839682 B1

TITLE: Predictive modeling of consumer financial behavior using supervised segmentation and nearest-neighbor matchingAbstract Text (1):

Predictive modeling of consumer financial behavior, including determination of likely responses to particular marketing efforts, is provided by application of consumer transaction data to predictive models associated with merchant segments. The merchant segments are derived from the consumer transaction data based on co-occurrences of merchants in sequences of transactions. Merchant vectors represent specific merchants, and are aligned in a vector space as a function of the degree to which the merchants co-occur more or less frequently than expected. Consumer vectors are developed within the vector space, to represent interests of particular consumers by virtue of relative vector positions of consumer and merchant vectors. Various techniques, including clustering, supervised segmentation, and nearest-neighbor analysis, are applied separately or in combination to generate improved predictions of consumer behavior.

Parent Case Text (2):

This application is a continuation-in-part of U.S. patent application Ser. No. 09/306,237 for "Predictive Modeling of Consumer Financial Behavior," filed May 6, 1999, now U.S. Pat. No. 6,430,539 the disclosure of which is incorporated by reference.

Brief Summary Text (3):

The present invention relates generally to analysis of consumer financial behavior, and more particularly to analyzing historical consumer financial behavior to accurately predict future spending behavior and likely responses to particular marketing efforts, in specifically identified data-driven industry segments.

Brief Summary Text (6):

Conventional means of determining consumer interests have generally relied on collecting demographic information about consumers, such as income, age, place of residence, occupation, and so forth, and associating various demographic categories with various categories of interests and merchants. Interest information may be collected from surveys, publication subscription lists, product warranty cards, and myriad other sources. Complex data processing is then applied to the source of data resulting in some demographic and interest description of each of a number of consumers.

Brief Summary Text (7):

This approach to understanding consumer behavior often misses the mark. The ultimate goal of this type of approach, whether acknowledged or not, is to predict consumer spending in the future. The assumption is that consumers will spend money on their interests, as expressed by things like their subscription lists and their demographics. Yet, the data on which the determination of interests is made is typically only indirectly related to the actual spending patterns of the consumer. For example, most publications have developed demographic models of their readership, and offer their subscription lists for sale to others interested in the particular demographics of the publication's readers. But subscription to a particular publication is a relatively poor indicator of what the consumer's spending patterns will be in the future.

Brief Summary Text (10):

Yet another problem with conventional approaches is that categorization of purchases is often based on standardized industry classifications of merchants and business, such as the SIC codes. This set of classification is entirely arbitrary, and has little to do with actual consumer behavior. Consumers do not decide which merchants to purchase from based on merchant

SIC codes. Thus, the use of arbitrary classifications to predict financial behavior is doomed to failure, since the classifications have little meaning in the actual data of consumer spending.

Brief Summary Text (13):

Accordingly, what is needed is the ability to model consumer financial behavior based on actual historical spending patterns that reflect the time-related nature of each consumer's purchase. Further, it is desirable to extract meaningful classifications of merchants based on the actual spending patterns, and from the combination of these, predict future spending of an individual consumer in specific, meaningful merchant groupings. Finally, it is desirable to provide recommendations based on analysis of customers that are similar to the target customer, and in particular to take into account the observed degree of success of particular marketing efforts with respect to such similar customers.

Brief Summary Text (14):

In the application domain of information, and particularly text retrieval, vector based representations of documents and words is known. Vector space representations of documents are described in U.S. Pat. No. 5,619,709 issued to Caid et. al, and in U.S. Pat. No. 5,325,298 issued to Gallant. Generally, vectors are used to represent words or documents. The relationships between words and between documents is learned and encoded in the vectors by a learning law. However, because these uses of vector space representations, including the context vectors of Caid, are designed for primarily for information retrieval, they are not effective for predictive analysis of behavior when applied to documents such as credit card statements and the like. When the techniques of Caid were applied to the prediction problems, it had numerous shortcomings. First, it had problems dealing with high transaction count merchants. These are merchants whose names appear very frequently in the collections of transaction statements. Because Caid's system downplays the significance of frequently appearing terms, these high transaction frequency merchants were not being accurately represented. Excluding high transaction frequency merchants from the data set however undermines the system's ability to predict transactions in these important merchants. Second, it was discovered that past two iterations of training, Caid's system performance declined, instead of converging. This indicates that the learning law is learning information that is only coincidental to transaction prediction, instead of information that is specifically for transaction prediction. Accordingly, it is desirable to provide a new methodology for learning the relationships between merchants and consumers so as to properly reflect the significance of the frequency with which merchants appears in the transaction data.

Brief Summary Text (16):

The present invention overcomes the limitations of conventional approaches to consumer analysis by providing a system and method of analyzing and predicting consumer financial behavior that uses historical, and time-sensitive, spending patterns of individual consumers. In one aspect, the invention generates groupings (segments) of merchants, which accurately reflect underlying consumer interests, and a predictive model of consumer spending patterns for each of the merchant segments. In another aspect, a supervised segmentation technique is employed to develop merchant segments that are of interest to the user. In yet another aspect, a "nearest neighbor" technique is employed, so as to identify those customers that are most similar to the target customer and to make predictions regarding the target customer based on observed behavior of the nearest neighbors. Current spending data of an individual consumer or groups of consumers can then be applied to the predictive models to predict future spending of the consumers in each of the merchant clusters, and/or marketing success data with respect to nearest neighbors can be applied to predict likelihood of success in promoting particular products to particular customers.

Brief Summary Text (17):

In one aspect, the present invention includes the creation of data-driven grouping of merchants, based essentially on the actual spending patterns of a group of consumers. Spending data of each consumer is obtained, which describes the spending patterns of the consumers in a time-related fashion. For example, credit card data demonstrates not merely the merchants and amounts spent, but also the sequence in which purchases were made. One of the features of the invention is its ability to use the co-occurrence of purchases at different merchants to group merchants into meaningful merchant segments. That is, merchants that are frequently shopped at within some number of transactions or time period of each other reflect a meaningful cluster. This data-driven clustering of merchants more accurately describes the interests or preferences

of consumers.

Brief Summary Text (18):

Merchants may also be segmented according to a supervised segmentation technique, such as Kohonen's Learning Vector Quantization (LVQ) algorithm, as described in T. Kohonen, "Improved Versions of Learning Vector Quantization," in IJCNN San Diego, 1990; and T. Kohonen, Self-Organizing Maps, 2d ed., Springer-Verlag, 1997. Supervised learning allows characteristics of segments to be directly specified, so that segments may be defined, for example, as "art museums," "book stores," "Internet merchants," and the like. Segment boundaries can be defined by the training algorithm based on training exemplars with known membership in classes. Segments may be overlapping or mutually exclusive, as desired.

Brief Summary Text (21):

In one embodiment, clustering techniques or supervised segmentation techniques are then applied to define merchant segments. Each merchant segment yields useful information about the type of merchants associated with it, their average purchase and transaction rates, and other statistical information. (Merchant "segments" and merchant "clusters" are used interchangeably herein.)

Brief Summary Text (23):

Preferably, each consumer is also given a profile that includes various demographic data, and summary data on spending habits. In addition, each consumer is preferably given a consumer vector. From the spending data, the merchants from whom the consumer has most frequently or recently purchased are determined. The consumer vector is then the summation of these merchant vectors. As new purchases are made, the consumer vector is updated, preferably decaying the influence of older purchases. In essence, like the expression "you are what you eat," the present invention reveals, "you are whom you shop at," since the vectors of the merchants are used to construct the vectors of the consumers.

Brief Summary Text (25):

Given the merchant segments, the present invention then creates a predictive model of future spending in each merchant segment, based on transaction statistics of historical spending in the merchant segment by those consumers who have purchased from merchants in the segments, in other segments, and data on overall purchases. In one embodiment, each predictive model predicts spending in a merchant cluster in a predicted time interval, such as 3 months, based on historical spending in the cluster in a prior time interval, such as the previous 6 months. During model training, the historical transactions in the merchant cluster for consumers who spent in the cluster, is summarized in each consumer's profile in summary statistics, and input into the predictive model along with actual spending in a predicted time interval. Validation of the predicted spending with actual spending is used to confirm model performance. The predictive models may be a neural network, or other multivariate statistical model.

Brief Summary Text (27):

To predict financial behavior, the consumer profile of a consumer, using preferably the same type of summary statistics for a recent, past time period, is input into the predictive models for the different merchant clusters. The result is a prediction of the amount of money that the consumer is likely to spend in each merchant cluster in a future time interval, for which no actual spending data may yet be available.

Brief Summary Text (28):

For each consumer, a membership function may be defined which describes how strongly the consumer is associated with each merchant segment. (Preferably, the membership function outputs a membership value for each merchant segment.) The membership function may be the predicted future spending in each merchant segment, or it may be a function of the consumer vector for the consumer and a merchant segment vector (e.g. centroid of each merchant segment). The membership function can be weighted by the amount spent by the consumer in each merchant segment, or other factors. Given the membership function, the merchant clusters for which the consumer has the highest membership values are of particular interest: they are the clusters in which the consumer will spend the most money in the future, or whose spending habits are most similar to the merchants in the cluster. This allows very specific and accurate targeting of promotions, advertising and the like to these consumers. A financial institution using the predicted spending information can direct promotional offers to consumers who are predicted to spend heavily in a merchant segment, with the promotional offers associated with merchants in

the merchant segment.

Brief Summary Text (29):

Also, given the membership values, changes in the membership values can be readily determined over time, to identify transitions by the consumer between merchants segments of interest. For example, each month (e.g. after a new credit card billing period or bank statement), the membership function is determined for a consumer, resulting in a new membership value for each merchant cluster. The new membership values can be compared with the previous month's membership values to indicate the largest positive and negative increases, revealing the consumer's changing purchasing habits. Positive changes reflect purchasing interests in new merchant clusters; negative changes reflect the consumer's lack of interest in a merchant cluster in the past month. Segment transitions such as these further enable a financial institution to target consumers with promotions for merchants in the segments in which the consumers show significant increases in membership values.

Brief Summary Text (30):

In another aspect, the present invention provides an improved methodology for learning the relationships between merchants in transaction data, and defining vectors that represent the merchants. More particularly, this aspect of the invention accurately identifies and captures the patterns of spending behavior that result in the co-occurrence of transactions at different merchants. The methodology is generally as follows:

Brief Summary Text (31):

First, the number of times that each pair of merchants co-occurs with one another in the transaction data is determined. The underlying intuition here is that merchants whom the consumers' behaviors indicates as being related will occur together often, whereas unrelated merchants do not occur together often. For example, a new mother will likely shop at children's clothes stores, toy stores, and other similar merchants, whereas a single young male will likely not shop at these types of merchants. The identification of merchants is by counting occurrences of merchants' names in the transaction data. The merchants' names may be normalized to reduce variations and equate different versions of a merchant's name to a single common name.

Brief Summary Text (39):

The present invention may be embodied in various forms. As a computer program product, the present invention includes a data preprocessing module that takes consumer spending data and processes it into organized files of account related and time organized purchases. Processing of merchant names in the spending data is provided to normalize variant names of individual merchants. A data post-processing module generates consumer profiles of summary statistics in selected time intervals, for use in training the predictive model. A predictive model generation system creates merchant vectors, and clusters them into merchant clusters, and trains the predictive model of each merchant segment using the consumer profiles and transaction data. Merchant vectors and consumer profiles are stored in databases. A profiling engine applies consumer profiles and consumer transaction data to the predictive models to provide predicted spending in each merchant segment, and to compute membership functions of the consumers for the merchant segment. A reporting engine outputs reports in various formats regarding the predicted spending and membership information. A segment transition detection engine computes changes in each consumer's membership values to identify significant transitions of the consumer between merchant clusters. The present invention may also be embodied as a system, with the above program product element cooperating with computer hardware components, and as a computer-implemented method.

Drawing Description Text (3):

FIG. 2 is a sample list of merchant segments.

Drawing Description Text (6):

FIG. 4b is an illustration of the system architecture of the present invention during development and training of merchant vectors, and merchant segment predictive models.

Drawing Description Text (11):

FIG. 9 is an illustration of the application of multiple consumer account data to the multiple segment predictive models.

Drawing Description Text (13):

FIGS. 11A through 11C show an example of segment vector adjustment.

Drawing Description Text (14):

FIGS. 12A through 12C show a second example of segment vector adjustment.

Detailed Description Text (3):

One feature of the present invention that enables prediction of consumer spending levels at specific merchants and prediction of response rates to marketing offers is the ability to represent both consumer and merchants in the same modeling representation. A conventional example is attempting to classify both consumers and merchants with demographic labels (e.g. "baby boomers", or "empty-nesters"). This conventional approach is simply arbitrary, and does not provide any mechanisms for directly quantifying how similar a consumer is to various merchants. The present invention, however, does provide such a quantifiable analysis, based on high-dimensional vector representations of both consumers and merchants, and the co-occurrence of merchants in the spending data of individual consumers.

Detailed Description Text (8):

Thus, in FIG. 1b, following processing of the consumer transaction data, the merchant vectors for merchants A, C, and E have been updated, based on actual spending data, such as C1's transactions, to point generally in the same direction, as have the merchant vectors for merchants B and D, based on C2's transactions. Clustering techniques are used then to identify clusters or segments of merchants based on their merchant vectors 402. In the example of FIG. 1b, a merchant segment is defined to include merchants A, C, and E, such as "upscale-technology_savvy." Note that as defined above, the SIC codes of these merchants are entirely unrelated, and so SIC code analysis would not reveal this group of merchants. Further, a different segment with merchants B and D is identified, even though the merchants share the same SIC codes with the merchants in the first segment, as shown in the transaction data 104.

Detailed Description Text (9):

Each merchant segment is associated with a merchant segment vector 105, preferably the centroid of the merchant cluster. Based on the types of merchants in the merchant segment, and the consumers who have purchased in the segment, a segment name can be defined, and may express the industry, sub-industry, geography, and/or consumer demographics.

Detailed Description Text (10):

The merchant segments provide very useful information about the consumers. In FIG. 1b there is shown the consumer vectors 106 for consumers C1 and C2. Each consumer's vector is a summary vector of the merchants at which the consumer shops. This summary is preferably the vector sum of merchant vectors at which the consumer has shopped at in defined recent time interval. The vector sum can be weighted by the recency of the purchases, their dollar amount, or other factors.

Detailed Description Text (11):

Being in the same vector space as the merchant vectors, the consumer vectors 106 reveal the consumer's interests in terms of their actual spending behavior. This information is by far a better base upon which to predict consumer spending at merchants, and likely response rates to offers, than superficial demographic labels or categories. Thus, consumer C1's vector is very strongly aligned with the merchant vectors of merchants A, C, and E, indicating C1 is likely to be interested in the products and services of these merchants. C1's vector can be aligned with these merchants, even if C1 never purchased at any of them before. Thus, merchants A, C, and E have a clear means for identifying consumers who may be interested in purchasing from them.

Detailed Description Text (12):

Which consumers are associated with which merchant segments can also be determined by a membership function. This function can be based entirely on the merchant segment vectors and the consumer vectors (e.g. dot product), or on other quantifiable data, such as amount spent by a consumer in each merchant segment, or a predicted amount to be spent.

Detailed Description Text (13):

Given the consumers who are members of a segment, useful statistics can be generated for the segment, such as average amount spent, spending rate, ratios of how much these consumers spend in the segment compared with the population average, response rates to offers, and so forth.

This information enables merchants to finely target and promote their products to the appropriate consumers.

Detailed Description Text (14):

FIG. 2 illustrates portions of a sample index of merchant segments, as may be produced by the present invention. Segments are named by assigning each segment a unique segment number 200 between 1 and M the total number of segments. In addition, each segment has a description field 210, which describes the merchant segment. A preferred description field is of the form:

Detailed Description Text (15):

Major Categories: Minor Categories: Demographics: Geography

Detailed Description Text (16):

Major categories 202 describe how the customers in a merchant segment typically use their accounts. Uses include retail purchases, direct marketing purchases, and where this type cannot be determined, then other major categories, such as travel uses, educational uses, services, and the like. Minor categories 204 describe both a subtype of the major category (e.g. subscriptions being a subtype of direct marketing) or the products or services purchased in the transactions (e.g. housewares, sporting goods, furniture) commonly purchased in the segment. Demographics information 206 uses account data from the consumers who frequent this segment to describe the most frequent or average demographic features, such as age range or gender, of the consumers. Geographic information 208 uses the account data to describe the most common geographic location of transactions in the segment. In each portion of the segment description 210 one or more descriptors may be used (i.e. multiple major, minor, demographic, or geographic descriptors). This naming convention is much more powerful and fine-grained than conventional SIC classifications, and provides insights into not just the industries of different merchants (as in SIC) but more importantly, into the geographic, approximate age or gender, and lifestyle choices of consumers in each segment.

Detailed Description Text (17):

The various types of segment reports are further described in section I. Reporting Engine, below.

Detailed Description Text (19):

Turning now to FIG. 4a there is shown an illustration of a system architecture of one embodiment of the present invention during operation in a mode for predicting consumer spending. System 400 includes begins with a data preprocessing module 402, a data postprocessing module 410, a profiling engine 412, and a reporting engine 426. Optional elements include a segment transition detection engine 420 and a targeting engine 422. System 400 operates on different types of data as inputs, including consumer summary file 404 and consumer transaction file 406, generates interim models and data, including the consumer profiles in profile database 414, merchant vectors 416, merchant segment predictive models 418, and produces various useful outputs including various segment reports 428-432.

Detailed Description Text (25):

The merchant vectors are then clustered 304 into merchant segments. The merchant segments generally describe groups of merchants that are naturally (in the data) shopped at "together" based on the transactions of the many consumers. Each merchant segment has a segment vector computed for it, which is a summary (e.g. centroid) of the merchant vectors in the merchant segment. Merchant segments provide very rich information about the merchants that are members of the segments, including statistics on rates and volumes of transactions, purchases, and the like.

Detailed Description Text (26):

With the merchant segments now defined, a predictive model of spending behavior is created 306 for each merchant segment. The predictive model for each segment is derived from observations of consumer transactions in two time periods: an input time window and a subsequent prediction time window. Data from transactions in the input time window for each consumer (including both segment specific and cross-segment) is used to extract independent variables, and actual spending in the prediction window provides the dependent variable. The independent variables typically describe the rate, frequency, and monetary amounts of spending in all segments and in the segment being modeled. A consumer vector derived from the consumer's transactions may also be used. Validation and analysis of the segment predictive models may be done to confirm the

Record Display Form

performance of the models.

Detailed Description Text (27):

In one embodiment, a predictive model may also be developed to predict spending at vendors, responses to particular offers or other marketing schemes, and the like, that are not associated with a particular market segment. The predictive model is trained using vector values of a number of customers with respect to a number of market segments. The customers' known spending behavior and/or responses to offers (both positive and negative exemplars) are provided as training data for the predictive model. Based on these data items, the model is trained, using known techniques such as neural network backward propagation techniques, linear regression, and the like. A predicted response or spending behavior estimate can then be generated based on vector values for a customer with respect to a number of market segments, even when the behavior being predicted does not correspond to any of the known market segments.

Detailed Description Text (28):

In the production phase, the system is used to predict spending, either in future time periods for which there is no actual data as of yet, or in a recent past time period for which data is available and which is used for retrospective analysis. Generally, each account (or consumer) has a profile summarizing the transactional behavior of the account holder. This information is created, or updated 308 with recent transaction data if present, to generate the appropriate variables for input into the predictive models for the segments. (Generation of the independent variables for model generation may also involve updating 308 of account profiles.)

Detailed Description Text (29):

Each account further includes a consumer vector which is derived, e.g. as a summary vector, from the merchant vectors of the merchant at which the consumer has purchased in a defined time period, say the last three months. Each merchant vector's contribution to the consumer vector can be weighted by the consumer's transactions at the merchants, such as by transaction amounts, rates, or recency. The consumer vectors, in conjunction with the merchant segment vectors provide an initial level of predictive power. Each consumer can now be associated with the merchant segment having a merchant segment vector closest to the consumer vector for the consumer.

Detailed Description Text (30):

Using the updated account profiles, this data is input into the set of predictive models to generate 310 for each consumer, an amount of predicted spending in each merchant segment in a desired prediction time period. For example, the predictive models may be trained on a six-month input window to predict spending in a subsequent three-month prediction window. The predicted period may be an actual future period or a current (e.g. recently ended) period for which actual spending is available.

Detailed Description Text (31):

The predicted spending levels and consumer profiles allow for various levels and types of account and segment analysis 312. First, each account may be analyzed to determine which segment (or segments) the account is a member of, based on various membership functions. A preferred membership function is the predicted spending value, so that each consumer is a member of the segment for which they have the highest predicted spending. Other measures of association between accounts and segments may be based on percentile rankings of each consumer's predicted spending across the various merchant segments. With any of these (or similar) methods of determining which consumers are associated with which segments, an analysis of the rates and volumes of different types of transactions by consumers in each segment can be generated. Further, targeting of accounts in one or more segments may be used to selectively identify populations of consumers with predicted high dollar amount or transaction rates. Account analysis also identifies consumers who have transitioned between segments as indicated by increased or decreased membership values.

Detailed Description Text (32):

Using targeting criteria, promotions directed 314 to specific consumers in specific segments and the merchants in those segments can be realized. For example, given a merchant segment, the consumers with the highest levels (or rankings) of predicted spending in the segment may be identified, or the consumers having consumer vectors closest to the segment vector may be selected. Or, the consumers who have highest levels of increased membership in a segment may be

selected. The merchants that make up the segment are known from the segment clustering 304. One or more promotional offers specific to merchants in the segment can be created, such as discounts, incentives and the like. The merchant-specific promotional offers are then directed to the selected consumers. Since these account holders have been identified as having the greatest likelihood of spending in the segment, the promotional offers beneficially coincide with their predicted spending behavior. This desirably results in an increased success rate at which the promotional offers are redeemed.

Detailed Description Text (33):

In an alternative embodiment, supervised segmentation is performed in place of the data-driven segmentation approach described above. Supervised segmentation allows a user to specify particular merchant segments that are of interest, so that relevant data can be extracted in a relevant and usable form. Examples of user-defined merchant segments include "art museums," "book stores," and "Internet merchants." Supervised segmentation allows a user to direct the system to provide predictive and analytical data concerning those particular segments in which the user is interested.

Detailed Description Text (34):

The technique of supervised segmentation, as employed by one embodiment of the present invention, determines segment boundaries and segment membership for merchants. Segment vectors are initialized, and are then iteratively adjusted using a training algorithm, until the segment vectors represent a meaningful summary of merchants belonging to the corresponding segment. The basis for the training algorithm is a Learning Vector Quantization (LVQ) technique, as described for example, in T. Kohonen, "Improved Versions of Learning Vector Quantization," in IJCNN San Diego, 1990. According to the techniques of the system, segments may overlap or they may be mutually exclusive, depending on user preference and the particular application. For example, with overlapping segments, a particular merchant (such as an Internet bookstore) might be a member of two or more merchant segments (e.g. "book stores" and "Internet merchants"). If mutually exclusive segments are used, the merchant will be assigned to only one segment, based on the learning algorithm's determination as to which segment is most suitable for the merchant.

Detailed Description Text (35):

Referring now to FIG. 10, there is shown a flowchart of an example of a supervised segmentation technique as may be used in connection with the present invention. According to the flowchart of FIG. 10, the system accepts user input specifying segments, and further specifying segment labels for a subset of merchants. Segment vectors are then iteratively adjusted based on the assigned segment labels, until segment vectors accurately represent an aggregation of the members of the respective segments.

Detailed Description Text (36):

A user specifies 1001 a set of merchant segments. A set of segment vectors are initialized 1002 for the specified merchant segments. The initial segment vectors may be orthogonal to one another, for example by being randomly assigned. Typically, the segment vectors occupy the same space as do merchant vectors, so that memberships, degrees of similarity, and affinities between merchants and segments can be defined and quantified.

Detailed Description Text (37):

For at least a subset of merchants, the user provides 1003 segment labels. In other words, the user labels the merchant with one (or more) of the specified merchant segments. These manually applied segment labels are then used by the system to train and refine segment vectors, as follows.

Detailed Description Text (38):

A labeled merchant is selected 1004 for processing. Based on the merchant vector for the selected merchant (derived previously from step 302 of FIG. 3, as described above), a segment is determined 1005 for the merchant. In one embodiment, this is the segment having a segment vector that is most closely aligned with the merchant vector (this may be determined, for example, by calculating the dot-product of the segment vector and merchant vector).

Detailed Description Text (39):

The segment specified by the manually applied segment label is compared 1006 with the segment determined in step 1005. If these are not the same segment, one or more segment vectors are

adjusted 1008 in an effort to "train" the segment vectors. Either the segment vector determined in 1006 is moved farther from the merchant vector, or the segment vector specified by the label is moved closer to the merchant vector, or both vectors are adjusted.

Detailed Description Text (40):

For example, suppose we have a merchant such as Barnes & Noble. The user provides a segment label identifying the merchant as a bookstore. In step 1005, the system determines a segment for the merchant based on vector positioning. If the determined segment is, for example, "grocery store," which does not match the segment label, segment vectors would be adjusted accordingly. The segment vector for grocery stores might be moved farther away from the Barnes & Noble merchant vector, or the segment vector for bookstores might be moved closer to the Barnes & Noble merchant vector, or both adjustments might be made.

Detailed Description Text (41):

Referring to FIGS. 11A through 11C, there are shown examples of segment vector adjustments that may be performed when the selected segment does not correspond to the segment label manually applied to the merchant vector. FIG. 11A depicts a starting position for a merchant vector MV and three segment vectors SV.sub.1, SV.sub.2, and SV.sub.3. For illustrative purposes, vector space 100 is depicted as having three dimensions, though in practice it is a hypersphere having any number of dimensions. MV is assumed to have been manually labeled with segment 1, corresponding to segment vector SV.sub.1. It can be seen from the starting positions shown in FIG. 11A that the segment vector closest to merchant vector MV is SV.sub.2, which does not correspond to the label assigned to MV. Accordingly, one or more of segment vectors SV.sub.1 and SV.sub.2 are adjusted.

Detailed Description Text (42):

FIG. 11B depicts an adjustment that may be performed on the segment vector SV.sub.2 that is closest to the merchant vector MV. Segment vector SV.sub.2 is moved away from MV, so as to reflect the fact that MV was not labeled with SV.sub.2 FIG. 11C depicts another adjustment that may be performed; in this figure, segment vector SV.sub.1 is moved closer to MV, so as to reflect the fact that MV was labeled with SV.sub.1. In an alternative embodiment, both adjustments depicted in FIGS. 11B and 11C may be performed.

Detailed Description Text (47):

If, in 1006, the selected segment does correspond to the segment label that has been assigned to the merchant, zero or more segment vectors are adjusted 1010. Either the segment vectors are left unchanged, or in an alternative embodiment, the assigned segment vector is moved closer to the merchant vector.

Detailed Description Text (48):

Thus continuing the Barnes & Noble example, if the determined segment is "bookstore," which does match the segment label, segment vectors may be left unchanged, or the segment vector for bookstores might be moved closer to the Barnes & Noble merchant vector.

Detailed Description Text (49):

Referring to FIGS. 12A through 12C, there is shown an example of a segment vector adjustment that may be performed when the selected segment does correspond to the segment label manually applied to the merchant vector. FIG. 12A depicts a starting position for a merchant vector MV and three segment vectors SV.sub.1, SV.sub.2, and SV.sub.3. MV is assumed to have been manually labeled with segment 1, corresponding to segment vector SV.sub.1. It can be seen from the starting positions shown in FIG. 12A that the segment vector closest to merchant vector MV is SV.sub.1, which does correspond to the label assigned to MV. Accordingly, either the vectors are left unchanged as shown in FIG. 12B, or, as shown in FIG. 12C, segment vector SV.sub.1 is moved closer to MV, so as to reflect the fact that MV was correctly assigned to SV.sub.1.

Detailed Description Text (52):

In yet another embodiment, segment membership is nonexclusive, so that a merchant may be a member of more than one segment. A tolerance radius is established around the endpoint of each segment vector along the surface of a unit sphere; this tolerance radius represents a maximum allowable distance from the vector endpoint to the endpoint of a merchant vector, along the surface of the sphere. The tolerance radius may also be expressed as a minimum value resulting from a dot-product operation on the segment vector and a merchant vector; if the dot-product value exceeds this threshold value, the merchant is designated a member of the segment. Either

technique may be used, as can any other method of defining a threshold value for segment membership.

Detailed Description Text (53):

Rather than adjusting segment vectors based on a determination of which segment vector is closest to the merchant vector, in this embodiment segment vectors are adjusted based on a determination of the merchant vector falling within the tolerance radius for one or more segment vectors. Adjustment of segment vectors may be performed as follows. Segment labels are manually applied to a merchant as described above in step 1003 of FIG. 10. The merchant vector is compared with segment vectors in order to determine whether the merchant vector falls within the predefined tolerance radius for any segment vectors. For each segment for which the merchant vector falls within the tolerance radius of the segment vector:

Detailed Description Text (54):

If the segment is one whose label was not manually applied to the merchant, adjust the segment vector to be farther from the merchant vector (FIG. 11B) and/or adjust other segment vectors whose labels were manually applied to the merchant to be closer to the merchant vector (FIG. 11C).

Detailed Description Text (55):

If the segment is one whose label was manually applied to the merchant, either do nothing (FIG. 12B) or adjust the segment vector to be closer to the merchant vector (FIG. 12C).

Detailed Description Text (56):

If the merchant vector does not fall within the tolerance radius of any segment vector, the system adjusts the segment vectors whose labels were manually applied to the merchant, to be closer to the merchant vector.

Detailed Description Text (57):

Once segments have been adjusted (if appropriate), a determination is made 1007 as to whether more training is required. This determination is made based on known convergence determination methods, or by reference to a predefined count of training iterations, or by other appropriate means. One advantage to the present invention is that not all merchants need be manually labeled in order to effectively train the vector set; once the segment vectors are sufficiently trained, merchants will automatically become associated with appropriate segments based on the positioning of their vectors.

Detailed Description Text (58):

As will be apparent to one skilled in the art, the supervised segmentation approach provides an alternative to unsupervised data-driven segmentation methods, and facilitates analysis of particular market segments or merchant types that are of interest. Thus, the above-described approach may be employed in place of the clustering methods previously described.

Detailed Description Text (61):

The data preprocessing module 402 (DPM) does initial processing of consumer data received from a source of consumer accounts and transactions, such as a credit card issuer, in preparation for creating the merchant vectors, consumer vectors, and merchant segment predictive models. DPM 402 is used in both production and training modes. (In this disclosure, the terms "consumer," "customer," and "account holder" are used interchangeably).

Detailed Description Text (63):

Customer summary file 404: The customer summary file 404 contains one record for each customer that is profiled by the system, and includes account information of the customer's account, and optionally includes demographic information about the customer. The consumer summary file 404 is typically one that a financial institution, such as a bank, credit card issuer, department store, and the like maintains on each consumer. The customer or the financial institution may supply the additional demographic fields that are deemed to be of informational or of predictive value. Examples of demographic fields include age, gender and income; other demographic fields may be provided, as desired by the financial institution.

Detailed Description Text (64):

Table 1 describes one set of fields for the customer summary file 404 for a preferred embodiment. Most fields are self-explanatory. The only required field is an account identifier

that uniquely identifies each consumer account and transactions. This account identifier may be the same as the consumer's account number; however, it is preferable to have a different identifier used, since a consumer may have multiple account relationships with the financial institution (e.g. multiple credit cards or bank accounts), and all transactions of the consumer should be dealt with together. The account identifier is preferably derived from the account number, such as by a one-way hash or encrypted value, such that each account identifier is uniquely associated with an account number. The pop_id field is optionally used to segment the population of customers into arbitrary distinct populations as specified by the financial institution, for example by payment history, account type, geographic region, etc.

Detailed Description Text (65):

Note the additional, optional demographic fields for containing demographic information about each consumer. In addition to demographic information, various summary statistics of the consumer's account may be included. These include any of the following:

Detailed Description Text (72):

a) Verify minimum data requirements. The DPM 402 determines the number of data files it is handling (since there maybe many physical media sources), and the length of the files to determine the number of accounts and transactions. Preferably, a minimum of 12 months of transactions for a minimum of 2 million accounts is used to provide fully robust models of merchants and segments. However, there is no formal lower bound to the amount of data on which system 400 may operate.

Detailed Description Text (77):

Referring to FIG. 4b, the predictive model generation system 440 takes as its inputs the master file 408 and creates the consumer profiles and consumer vectors, the merchant vectors and merchant segments, and the segment predictive models. This data is used by the profiling engine to generate predictions of future spending by a consumer in each merchant segment using inputs from the data postprocessing module 410.

Detailed Description Text (163):

The second technique, UDL2, overcomes of the small count problem by using log-likelihood ratio estimates to calculate $r_{sub.ij}$. It has been shown that log-likelihood ratios have much better small count behavior than χ^2 , while at the same time retaining the same behavior as χ^2 in the non-small count regions.

Detailed Description Text (184):

Following generation and training of the merchant vectors, the clustering module 520 is used to cluster the resulting merchant vectors and identify the merchant segments. Various different clustering algorithms may be used, including k-means clustering (MacQueen). The output of the clustering is a set of merchant segment vectors, each being the centroid of a merchant segment, and a list of merchant vectors (thus merchants) included in the merchant segment.

Detailed Description Text (185):

There are two different clustering approaches that may be usefully employed to generate the merchant segments. First, clustering may be done on the merchant vectors themselves. This approach looks for merchants having merchant vectors which are substantially aligned in the vector space, and clusters these merchants into segments and computes a cluster vector for each segment. Thus, merchants for whom transactions frequently co-occur and have high dot products between their merchant vectors will tend to form merchant segments. Note that it is not necessary for all merchants in a cluster to all co-occur in many consumers' transactions. Instead, co-occurrence is associative: if merchants A and B co-occur frequently, and merchants B and C co-occur frequently, A and C are likely to be in the same merchant segment.

Detailed Description Text (190):

However computed, the consumer vectors can then be clustered, so that similar consumers, based on their purchasing behavior, form a merchant segment. This defines a merchant segment vector. The merchant vectors that are closest to a particular merchant segment vector are deemed to be included in the merchant segment.

Detailed Description Text (191):

With the merchant segments and their segment vectors, the predictive models for each segment may be developed. Before discussing the creation of the predictive models, a description of the

training data used in this process is described.

Detailed Description Text (193):

Following identification of merchant segments, a predictive model of consumer spending in each segment is generated from past transactions of consumers in the merchant segment. Using the past transactions of consumer in the merchant segment provides a robust base on which to predict future spending, and since the merchant segments were identified on the basis of the actual spending patterns of the consumers, the arbitrariness of conventional demographic based predictions are minimized. Additional non-segment specific transactions of the consumer may also be used to provide a base of transaction behavior.

Detailed Description Text (194):

To create the segment models, the consumer transaction data is organized into groups of observations. Each observation is associated with a selected end-date. The end-date divides the observation into a prediction window and an input window. The input window includes a set of transactions in a defined past time interval prior to the selected end-date (e.g. 6 months prior). The prediction window includes a set of transactions in a defined time interval after the selected end-date (e.g. the next 3 months). The prediction window transactions are the source of the dependent variables for the prediction, and the input window transactions are the source of the independent variables for the prediction.

Detailed Description Text (196):

The first type of observations is training observations, which are used to train the predictive model that predicts future spending within particular merchant segments. If N is the length (in months) of the window over which observation inputs are computed then there are $2N-1$ training observations for each segment.

Detailed Description Text (198):

The second type of observations is blind observations. Blind observations are observations where the prediction window does not overlap any of the time frames for the prediction windows in the training observations. Blind observations are used to evaluate segment model performance. In FIG. 8, the blind observations 804 include those from September to February, as illustrated.

Detailed Description Text (200):

FIG. 8 also illustrates that at some point during the prediction window, the financial institution sends out promotions to selected consumers based on their predicted spending in the various merchant segments.

Detailed Description Text (201):

Referring to FIG. 4b again, the DPPM takes the master files 408, and a given selected end-date, and constructs for each consumer, and then for each segment, a set of training observations and blind observations from the consumer's transactions, including transactions in the segment, and any other transactions. Thus, if there are 300 segments, for each consumer there will be 300 sets of observations. If the DPPM is being used during production for prediction purposes, then the set of observations is a set of action observations.

Detailed Description Text (203):

Prediction window: The dependent variables are generally any measure of amount or rate of spending by the consumer in the segment in the prediction window. A simple measure is the total dollar amount that was spent in the segment by the consumer in the transactions in the prediction window. Another measure may be average amount spent at merchants (e.g. total amount divided by number of transactions).

Detailed Description Text (204):

Input window: The independent variables are various measures of spending in the input window leading up to the end date (though some may be outside of it). Generally, the transaction statistics for a consumer can be extracted from various grouping of merchants. These groups may be defined as: 1) merchants in all segments; 2) merchants in the merchant segment being modeled; 3) merchants whose merchant vector is closest the segment vector for the segment being modeled (these merchants may or may not be in the segment); and 4) merchants whose merchant vector is closest to the consumer vector of the consumer.

Detailed Description Text (206):

(1) Recency. The amount of time in months between the current end date and the most recent transaction of the consumer in any segment. Recency may be computed over all available time and is not restricted to the input window.

Detailed Description Text (207):

(2) Frequency. The number of transactions by a consumer in the input window preceding the end-date for all segments.

Detailed Description Text (208):

(3) Monetary value of purchases. A measure of the amount of dollars spent by a customer in the input window preceding the end-date for all segments. The total or average, or other measures may be used.

Detailed Description Text (209):

(4) Recency segment. The amount of time in months between the current end date and the most recent transaction of the consumer in the segment. Recency may be computed over all available time and is not restricted to the input window.

Detailed Description Text (210):

(5) Frequency segment. The number of transactions in the segment by a customer in the input window preceding the current end date.

Detailed Description Text (211):

(6) Monetary segment. The amount of dollars spent in the segment by a customer in the input window preceding the current end date.

Detailed Description Text (212):

(7) Recency nearest profile merchants. The amount of time in months between the current end date and the most recent transaction of the consumer in a collection of merchants that are nearest the consumer vector of the consumer. Recency may be computed over all available time and is not restricted to the input window.

Detailed Description Text (213):

(8) Frequency nearest profile merchants. The number of transactions in a collection of merchants that are nearest the consumer vector of the consumer in the input window preceding the current end date.

Detailed Description Text (215):

(10) Recency nearest segment merchants. The amount of time in months between the current end date and the most recent transaction of the consumer in a collection of merchants that are nearest the segment vector. Recency may be computed over all available time and is not restricted to the input window.

Detailed Description Text (216):

(11) Frequency nearest segment merchants. The number of transactions in a collection of merchants that are nearest the segment vector by the consumer in the input window preceding the current end date.

Detailed Description Text (217):

(12) Monetary nearest segment merchants. The amount of dollars spent in a collection of merchants that are nearest the segment vector by the consumer in the input window preceding the current end date.

Detailed Description Text (218):

(13) Segment probability score. The probability that a consumer will spend in the segment in the prediction window given all merchant transactions for the consumer in the input window preceding the end date. A preferred algorithm estimates combined probability using a recursive Bayesian method.

Detailed Description Text (220):

(15) (Segment Vector-Consumer Vector Closeness: As an optional input, the dot product of the segment vector for the segment and the consumer vector is used as an input variable.

Detailed Description Text (221):

In addition to these transaction statistics, variables may be defined for the frequency of purchase and monetary value for all cases of segment merchants, nearest profile merchants, nearest segment merchants for the same forward prediction window in the previous year(s).

Detailed Description Text (223):

The training observations for each segment are input into the segment predictive model generation module 530 to generate a predictive model for the segment. FIG. 9 illustrates the overall logic of the predictive model generation process. The master files 408 are organized by accounts, based on account identifiers, here illustratively, accounts 1 through N. There are M segments, indicated by segments 1 through M. The DPPM generates for each combination of account and merchant segment, a set of input and blind observations. The respective observations for each merchant segment M from the many accounts 1 . . . N are input into the respective segment predictive model M during training. Once trained, each segment predictive model is tested with the corresponding blind observations. Testing may be done by comparing for each segment a lift chart generated by the training observations with the lift chart generated from blind observations. Lift charts are further explained below.

Detailed Description Text (227):

the profiling engine 412 provides analytical data in the form of an account profile about each customer whose data is processed by the system 400. The profiling engine is also responsible for updating consumer profiles over time as new transaction data for consumers is received. The account profiles are objects that can be stored in a database 414 and are used as input to the computational components of system 400 in order to predict future spending by the customer in the merchant segments. The profile database 414 is preferably ODBC compliant, thereby allowing the accounts provider (e.g. financial institution) to import the data to perform SQL queries on the customer profiles.

Detailed Description Text (228):

The account profile preferably includes a consumer vector, a membership vector describing a membership value for the consumer for each merchant segment, such as the consumer's predicted spending in each segment in a predetermined future time interval, and the recency, frequency, and monetary variables as previously described for predictive model training.

Detailed Description Text (229):

The profiling engine 412 creates the account profiles as follows.

Detailed Description Text (230):

1. Membership Function: Predicted Spending in Each Segment

Detailed Description Text (231):

The profile of each account holder includes a membership value with respect to each segment. The membership value is computed by a membership function. The purpose of the membership function is to identify the segments with which the consumer is mostly closely associated, that is, which best represent the group or groups of merchants at which the consumer has shopped, and is likely to shop at in the future.

Detailed Description Text (232):

In a preferred embodiment, the membership function computes the membership value for each segment as the predicted dollar amount that the account holder will purchase in the segment given previous purchase history. The dollar amount is projected for a predicted time interval (e.g. 3 months forward) based on a predetermined past time interval (e.g. 6 months of historical transactions). These two time intervals correspond to the time intervals of the input window and prediction windows used during training of the merchant segment predictive models. Thus, if there are 300 merchant segments, then a membership value set is a list of 300 predicted dollar amounts, corresponding to the respective merchant segments. Sorting the list by the membership value identifies the merchant segments at which the consumer is predicted to spend the greatest amounts of money in the future time interval, given their spending historically.

Detailed Description Text (233):

To obtain the predicted spending, certain data about each account is input in each of the

segment predictive models. The input variables are constructed for the profile consistent with the membership function of the profile. Preferably, the input variables are the same as those used during model training, as set forth above. An additional input variable for the membership function may include the dot product between the consumer vector and the segment vector for the segment (if the models are so trained). The output of the segment models is a predicted dollar amount that the consumer will spend in each segment in the prediction time interval.

Detailed Description Text (234):

2. Segment Membership Based on Consumer Vectors

Detailed Description Text (235):

A second alternate, membership aspect of the account profiles is membership based upon the consumer vector for each account profile. The consumer vector is a summary vector of the merchants that the account has shopped at, as explained above with respect to the discussion of clustering. In this aspect, the dot product of the consumer vector and segment vector for the segment defines a membership value. In this embodiment, the membership value list is a set of 300 dot products, and the consumer is member of the merchant segment(s) having the highest dot product(s).

Detailed Description Text (236):

With either one of these membership functions, the population of accounts that are members of each segment (based on the accounts having the highest membership values for each segment) can be determined. From this population, various summary statistics about the accounts can be generated such as cash advances, purchases, debits, and the like. This information is further described below.

Detailed Description Text (237):

3. Updating of Consumer Profiles

Detailed Description Text (247):

The reporting engine 426 provides various types of segment and account specific reports. The reports are generated by querying the profiling engine 412 and the account database for the segments and associated accounts, and tabulating various statistics on the segments and accounts.

Detailed Description Text (254):

2. General Segment Report

Detailed Description Text (255):

For each merchant segment a very detailed and powerful analysis of the segment can be created in a segment report. This information includes:

Detailed Description Text (256):

a) General Segment Information

Detailed Description Text (257):

Merchant Cohesion: A measure of how closely clustered are the merchant vectors in this segment. This is the average of the dot products of the merchant vectors with the centroid vector of this segment. Higher numbers indicate tighter clustering.

Detailed Description Text (258):

Number of Transactions: The number of purchase transactions at merchants in this segment, relative to the total number of purchase transactions in all segments, providing a measure of how significant the segment is in transaction volume.

Detailed Description Text (259):

Dollars Spent: The total dollar amount spent at merchants in this segment, relative to the total dollar amount spent in all segments, providing a measure of dollar volume for the segment.

Detailed Description Text (260):

Most Closely Related Segments: A list of other segments that are closest to the current segment. This list may be ranked by the dot products of the segment vectors, or by a measure of

the conditional probability of purchase in the other segment given a purchase in the current segment.

Detailed Description Text (261):

The conditional probability measure M is as follows: $P(A.\text{vertline}.B)$ is probability of purchase in segment A segment in next time interval (e.g. 3 months) given purchases in segment B in the previous time interval (e.g. 6 months). $P(A.\text{vertline}.B)/P(A)=M$. If M is >1 , then a purchase in segment B is positively influencing the probability of purchase in segment A, and if $M<1$ then a purchase in segment B negatively influences a purchase in segment A. This is because if there is no information about the probability of purchases in segment B, then $P(A.\text{vertline}.B)=P(A)$, so $M=1$. The values for $P(A.\text{vertline}.B)$ are determined from the co-occurrences of purchases at merchants in the two segments, and $P(A)$ is determined and from the relative frequency of purchases in segment A compared to all segments.

Detailed Description Text (262):

A farthest segments list may also be provided (e.g. with the lowest conditional probability measures).

Detailed Description Text (263):

b) Segment Members Information

Detailed Description Text (264):

Detailed information is provided about each merchant that is a member of a segment. This information comprises:

Detailed Description Text (266):

Dollar Bandwidth: The fraction of all the money spent in this segment that is spent at this merchant (percent);

Detailed Description Text (269):

Merchant Score: The dot product of this merchant's vector with the centroid vector of the merchant segment. (A value of 1.0 indicates that the merchant vector is at the centroid);

Detailed Description Text (274):

Tables 10 illustrates a sample lift chart for merchant segment:

Detailed Description Text (277):

For each merchant segment then, the consumer accounts are ranked by their predicted spending for the segment in the prediction window period. Once the accounts are ranked, they are divided into N (e.g. 20) equal sized bins so that bin 1 has the highest spending accounts, and bin N has the lowest ranking accounts. This identifies the accounts holders that the predictive model for the segment indicated should be are expected to spend the most in this segment.

Detailed Description Text (278):

Then, for each bin, the average actual spending per account in this segment in the past time period, and the average predicted spending is computed. The average actual spending over all bins is also computed. This average actual spending for all accounts is the baseline spending value (in dollars), as illustrated in the last line of Table 10. This number describes the average that all account holders spent in the segment in the prediction window period.

Detailed Description Text (279):

The lift for a bin is the average actual spending by accounts in the bin divided by the baseline spending value. If the predictive model for the segment is accurate, then those accounts in the highest ranked bins should have a lift greater than 1, and the lift should generally be increasing, with bin 1 having the highest lift. Where this the case, as for example, in Table 10, in bin 1, this shows that those accounts in bin 1 in fact spent several times the baseline, thereby confirming the prediction that these accounts would in fact spend more than others in this segment.

Detailed Description Text (281):

The lift information allows the financial institution to very selectively target a specific group of accounts (e.g. the accounts in bin 1) with promotional offers related to the merchants in the segment. This level of detailed, predictive analysis of very discrete groups of specific

accounts relative to merchant segments is not believed to be currently available by conventional methods.

Detailed Description Text (283):

The reporting engine 426 further provides two types of analyses of the financial behavior of a population of accounts that are associated with a segment based on various selection criteria. The Segment Predominant Scores Account Statistics table and the Segment Top 5% Scores Account Statistics table present averaged account statistics for two different types of populations of customers who shop, or are likely to shop, in a given segment. The two populations are determined as follows.

Detailed Description Text (284):

Segment Predominant Scores Account Statistics Table

Detailed Description Text (285):

All open accounts with at least one purchase transaction are scored (predicted spending) for all of the segments. Within each segment, the accounts are ranked by score, and assigned a percentile ranking. The result is that for each account there is a percentile ranking value for each of the merchant segments.

Detailed Description Text (286):

The population of interest for a given segment is defined as those accounts that have their highest percentile ranking in this segment. For example, if an account has its highest percentile ranking in segment #108, that account will be included in the population for the statistics table for segment #108, but not in any other segment. This approach assigns each account holder to one and only one segment.

Detailed Description Text (287):

Segment Top 5% Scores Account Statistics

Detailed Description Text (288):

For the Segment Top 5% Scores Account Statistics table, the population is defined as the accounts with percentile ranking of 95% or greater in a current segment. These are the 5% of the population that is predicted to spend the most in the segment in the predicted future time interval following the input data time window. These accounts may appear in this population in more than one segment, so that high spenders will show up in many segments; concomitantly, those who spend very little may not be assigned to any segment.

Detailed Description Text (291):

i) Segment Statistics

Detailed Description Text (295):

Population Mean: the average, over all the segments, of the Mean Value (this column is thus the same for all segments, and are included for ease of comparison); and

Detailed Description Text (302):

The Dollars in Segment shows the fraction of total spending that is spent in this segment. This informs the financial institution of how significant overall this segment is.

Detailed Description Text (303):

The Rate in Segment shows the fraction of total purchase transactions that occur in this segment.

Detailed Description Text (304):

The differences between these two populations are subtle but important, and are illustrated by the above tables. The segment predominant population identifies those individuals as members of a segment who, relative to their own spending, are predicted to spend the most in the segment. For example, assume a consumer whose predicted spending in a segment is \$20.00, which gives the consumer a percentile ranking of 75.sup.th percentile. If the consumer's percentile ranking in every other segment is below the 75.sup.th percentile, then the consumer is selected in this population for this segment. Thus, this may be considered an intra-account membership function.

Detailed Description Text (305):

The Top 5% scores population instead includes those accounts holders predicted to spend the most in the segment, relative to all other account holders. Thus, the account holder who was predicted to spend only \$20.00 in the merchant segment will not be member of this population since he is well below the 95.sup.th percentile, which may be predicted to spend, for example \$100.00.

Detailed Description Text (306):

In the example tables these differences are pronounced. In Table 11, the average purchases of the segment predominant population is only \$166.86. In Table 12, the average purchase by top 5% population is more than twice that, at \$391.54. This information allows the financial institution to accurately identify accounts that are most likely to spend in a given segment, and target these accounts with promotional offers for merchants in the segment.

Detailed Description Text (307):

The above tables may also be constructed based on other functions to identify accounts associated with segments, including dot products between consumer vectors and segment vectors.

Detailed Description Text (309):

The targeting engine 422 allows the financial institution to specify targeted populations for each (or any) merchant segment, to enable selection of the targeted population for receiving predetermined promotional offers.

Detailed Description Text (310):

A financial institution can specify a targeted population for a segment by specifying a population count for the segment, for example, the top 1000 accounts holders, or the top 10% account holders in a segment. The selection is made by any of the membership functions, including dot product, or predicted spending. Other targeting specifications may be used in conjunction with these criteria, such as a minimum spending amount in the segment, such as \$100. The parameters for selecting the targeting population are defined in a target specification document 424, which is an input to the targeting engine 422. One or more promotions can be specifically associated with certain merchants in a segment, such as the merchants with the highest correlation with the segment vector, highest average transaction amount, or other selective criteria. In addition, the amounts offered in the promotions can be specific to each consumer selected, and based on their predicted or historical spending in the segment. The amounts may also be dependent on the specific merchant for whom a promotion is offered, as a function of the merchant's contributions to purchases in the segment, such as based upon their dollar bandwidth, average transaction amount, or the like.

Detailed Description Text (311):

The selected accounts can be used to generate a targeted segmentation report 430 by providing the account identifiers for the selected accounts to the reporting engine 426, which constructs the appropriate targeting report on the segment. This report has the same format as the general segment report but is compiled for the selected population.

Detailed Description Text (313):

Table 13 shows a specification of a total of at least 228,000 customer accounts distributed over four segments and two promotional offers (ID 1 and ID 2). For each segment or promotional offer, there are different selection and filtering criteria. For promotion #1 the top 75,000 consumers in segment #122 based on predicted spending, and who have an average transaction in the segment greater than \$50, are selected. For this promotion in segment #413, the top 10% of accounts based on the dot product between the consumer vector and segment vector are selected, so long as they have a minimum spending in the segment of \$100. Finally, for promotion #2, 87,000 consumers are selected across two segments. Within each offer (e.g. offer ID 1) the segment models may be merged to produce a single lift chart, which reflects the offer as a composition of the segments.

Detailed Description Text (315):

1. Select fields from the account profile of the selected accounts that will be inserted to the mail file 434. For example, the name, address, and other information about the account may be extracted.

Detailed Description Text (318):

4. Instruct the reporting engine 426 to generate lift charts for the targeting population in the segment, and for overlapped (combined) segments.

Detailed Description Text (319):

The predictive model may also be trained to predict spending at vendors, responses to particular offers or other marketing schemes, and the like, that are not associated with a particular market segment. Referring now to FIG. 13, training set 1301 contains data describing customers who have previously been presented with the offer, including customers who accepted the offer (positive exemplars) and customers who rejected the offer (negative exemplars). Vector values in the appropriate merchant vector space are also provided. Based on the data in training set 1301, predictive model 1303 is trained using known techniques, such as those of predictive model generation module 530 as referenced above.

Detailed Description Text (320):

Once a trained model 1303 is available, predicted response 1304 for a customer can be generated based on vector values 1304 for the customer in a number of merchant segments. The particular response 1304 being predicted need not be associated with any particular market segment in order for an effective prediction to be generated. In this manner, the system is able to provide meaningful predictions even in industries or marketing environments where market segments are not available or are inapplicable.

Detailed Description Text (321):

For example, suppose a prediction is to be generated for a particular consumer's response to an offer for a home equity line of credit. Training set 1301 would include some aggregation of data that describes the responses to the same (or similar) offer of a number of consumers. Vector values for those consumers in a number of market segments, along with the responses to the offer, would be used to train predictive model 1303. Then, given the particular consumer's vector values for a number of market segments 1304, model 1303 is able to predict the consumer's response 1304 to the offer for the line of credit, even though no market segment has been established for the offer.

Detailed Description Text (322):

K. Segment Transition Detection

Detailed Description Text (323):

As is now apparent, the system of the present invention provides detailed insight into which merchant segments a consumer is associated with based on various measures of membership, such as dot product, predicted spending, and the like. Further, since the consumers continue to spend over time, the consumer accounts and the consumers' associations with segments are expected to change over time as their individual spending habits change.

Detailed Description Text (324):

The present invention allows for detection of the changes in consumer spending via the segment transition detection engine 420. In a given data period (e.g. next monthly cycle or multiple month collection of data) a set of membership values for each consumer is defined as variously described above, with respect to each segment. Again, this may be predicted spending by the consumer in each segment, dot product between the consumer vector and each segment vectors, or other membership functions.

Detailed Description Text (325):

In a subsequent time interval, using additional spending and/or predicted data, the membership values are recomputed. Each consumer will have the top P and the bottom Q increases in and decreases in segment membership. That is, there will be two changes of interest: the P (e.g. 5) segments with the greatest increase in membership values for the consumer; the Q segments with the greatest decrease in segment membership.

Detailed Description Text (326):

An increase in the membership value for a segment indicates that the consumer is now spending (or predicted to spend) more money in a particular segment. Decreases show a decline in the consumer's interest in the segment. Either of these movements may reflect a change in the consumer's lifestyle, income, or other demographic factors.

Detailed Description Text (327):

Significant increases in merchant segments that previously had low membership values are particularly useful to target promotional offers to the account holders who are moving into the segment. This is because the significant increase in membership indicates that the consumer is most likely to be currently receptive to the promotional offers for merchants in the segment, since they are predicted to be purchasing more heavily in the segment.

Detailed Description Text (328):

Thus, the segment transition detection engine 420 calculates the changes in each consumer's membership values between two selected time periods, typically using data in a most recent prediction window (either ending or beginning with a current statement date) relative to memberships in prior time intervals. The financial institution can define a threshold change value for selecting accounts with changes in membership more significant than the threshold. The selected accounts may then be provided to the reporting engine 426 for generation of various reports, including a segment transition report 432, which is like the general segment report except that it applies to accounts that are considered to have transitioned to or from a segment. This further enables the financial institution to selectively target these customers with promotional offers for merchants in the segments in which the consumer had the most significant positive increases in membership.

Detailed Description Text (341):

In one embodiment, the nearest-neighbor response rate may be fused with other data for more advanced analysis. For example, the aggregated response rate could be provided as an input to a second-level predictive model, along with other input data (such as demographic information, for example). The second-level predictive model could be trained on the input data, using techniques known in the art, in order to improve response prediction accuracy for target consumers. Thus, the second-level predictive model would learn relationships among aggregated response rates and other input data, in order to generate a second-level predicted response rate that yields improved results. The relationships are learned using conventional training techniques, such as backward propagation and the like.

Detailed Description Text (344):

In addition, random sampling tends to yield many more non-responders and negative responders than positive responders, by virtue of the fact that, in general, the vast majority of people respond negatively (or not at all) to offers. Thus, random selection of reference consumers tends to result in an undue emphasis on non-responders and negative responders, with a corresponding lack of predictive data points for positive responders. This is an unfavorable result, since it weakens the ability of the system to develop sufficient numbers of vectors for the very population segment that is of the most interest, namely those who responded positively in the past.

Detailed Description Text (347):

Supervised segmentation of merchant vectors is described above as a technique for developing merchant segments that are of interest. In one embodiment, the system employs supervised segmentation of consumer vectors as an alternative to the nearest-neighbor technique described above for predicting response rates of consumers. Such a technique may be performed, for example, using an LVQ methodology similar to that described above in connection with merchant vectors.

Detailed Description Text (348):

Referring now to FIG. 15, there is shown a flowchart depicting a technique of supervised segmentation of consumer vectors for predicting a response rate for a consumer with regard to a particular offer. A set of reference consumers is labeled 1501 according to their response history for an offer. For each product offer, there are two classes of individuals--responders (those who responded positively) and non-responders (those who responded negatively or did not respond at all). Alternatively, multiple segment vectors can be trained with different ratios (or ranges of ratios) of responders to non-responders in order to model the response likelihood contours in the feature space.

Detailed Description Text (349):

A set of segment vectors are initialized 1502 for the specified consumer segments. The initial segment vectors may be orthogonal to one another, for example by being randomly assigned.

Detailed Description Text (350):

Typically, the segment vectors occupy the same space as do consumer vectors, so that memberships, degrees of similarity, and affinities between consumers and segments can be defined and quantified. In another alternative embodiment, more than one segment vector may be assigned to each segment in order to identify discontinuous regions of high response likelihood and to better approximate the decision boundaries.

Detailed Description Text (351):

A labeled reference consumer is selected 1503. A consumer vector is obtained for the selected reference consumer, and a segment is selected 1504 for the consumer based on the consumer vector. As described previously, segment selection may be performed according to any one of several methods, including for example determining which segment vector is most closely aligned with the consumer vector. If, in 1505, the selected segment does not correspond to the segment label that has been assigned to the consumer, one or more segment vectors are adjusted 1506 in an effort to "train" the segment vectors. Either the segment vector for the assigned segment is moved farther from the consumer vector, or the "correct" segment vector (i.e., the segment vector closest to the consumer vector) is moved closer to the segment vector, or both vectors are adjusted. Examples discussed above in connection with FIGS. 11A through 11C and 12A through 12C are applicable.

Detailed Description Text (352):

Once segments have been adjusted (if appropriate), a determination is made 1507 as to whether more training is required. This determination is made based on known convergence determination methods, or by reference to a predefined count of training iterations, or by another other appropriate means. One advantage to the system of present invention is that not all consumer vectors need be manually labeled in order to effectively train the vector set; once the segment vectors are sufficiently trained, consumers will automatically become associated with appropriate segments based on the positioning of their vectors.

Detailed Description Text (353):

Thus, the system provides a technique for developing segment vectors such that probability of response for each region of feature space may be determined. For a new target customer, the consumer vector is compared with segment vectors; based on a determination of response rate for a corresponding segment vector, the estimated response probability for the target customer can be generated. Such a technique is advantageous in that it results in reduced search time over a nearest-neighbor technique, and is more likely to provide accurate results in the presence of abrupt response likelihood boundaries in the feature space.

Detailed Description Text (354):

In summary then, the present invention provides a variety of powerful analytical methods for predicting consumer financial behavior in discretely defined merchant segments, and with respect to predetermined time intervals. The clustering of merchants in merchant segments allows analysis of transactions of consumers in each specific segment, both historically, and in the predicted period to identify consumers of interest. Identified consumers can then be targeted with promotional offers precisely directed at merchants within specific segments. Supervised segmentation techniques may be employed to facilitate definition and analysis of particular market segments. Nearest-neighbor techniques may be used in place of segment-based models to develop predictions of consumer behavior.

Detailed Description Paragraph Table (1):

TABLE 1 Customer Summary File Description Sample Format Account_id Char[max 24] Pop_id Char ('1'-'N') Account_number Char[max 16] Credit bureau Short int as score string Internal credit risk Short int as score string Ytd purchases Int as string Ytd_cash_adv Int as string Ytd_int_purchases Int as string Ytd_int_cash_adv Int as string State_code Char[max 2] Zip_code Char[max 5] Demographic_1 Int as string . . . Demographic_N Int as string

Detailed Description Paragraph Table (2):

TABLE 2 Example Demographic Fields for Customer Summary File Description Explanation Cardholder zip code Months on books or open date Number of people on the Equivalent to number of plastics account Credit risk score Cycles delinquent Credit line Open to buy Initial month statement bal- Balance on the account prior to once the first month of transaction data pull Last month statement balance Balance on the account at the end of the transaction data pulled Monthly payment amount For each month of transaction data contributed or the average over last year. Monthly cash advance For each month of transaction amount data contributed or the average over

last year. Monthly cash advance count For each month of transaction data contributed or the average over last year. Monthly purchase amount For each month of transaction data contributed or the average over last year. Monthly purchase count For each month of transaction data contributed or the average over last year. Monthly cash advance interest For each month of transaction data contributed or the average over last year. Monthly purchase interest For each month of transaction data contributed or the average over last year. Monthly late charge For each month of transaction data contributed or the average over last year.

Detailed Description Paragraph Table (4):

TABLE 4 Master File 408 Description Sample Format Account_id Char[max 24] Pop_id Char ('1'-'N') Account_number Char[max 16] Credit bureau score Short int as string Ytd purchases Int as string Ytd_cash_advances Int as string Ytd_interest_on_purchases Int as string Ytd_interest_on_cash_a Int as string dvs State_code Char[max 2] Demographic_1 Int as string . . . Demographic_N Int as string <transactions>

Detailed Description Paragraph Table (10):

TABLE 10 A sample segment lift chart. Cumulative Cumulative Cumulative Bin segment lift segment lift in \$ Population 1 5.56 \$109.05 50,000 2 4.82 \$94.42 100,000 3 3.82 \$74.92 150,000 4 3.23 \$63.38 200,000 5 2.77 \$54.22 250,000 6 2.43 \$47.68 300,000 7 2.20 \$43.20 350,000 8 2.04 \$39.98 400,000 9 1.88 \$36.79 450,000 10 1.75 \$34.35 500,000 11 1.63 \$31.94 550,000 12 1.52 \$29.75 600,000 13 1.43 \$28.02 650,000 14 1.35 \$26.54 700,000 15 1.28 \$25.08 750,000 16 1.21 \$23.81 800,000 17 1.16 \$22.65 850,000 18 1.10 \$21.56 900,000 19 1.05 \$20.57 950,000 20 1.00 \$19.60 1,000,000 Base-line -- \$19.60

Detailed Description Paragraph Table (11):

TABLE 11 Segment Predominant Scores Account Statistics: 8291 accounts (0.17 per- cent) Population Relative Category Mean Value Std Deviation Mean Score Cash Advances \$11.28 \$53.18 \$6.65 169.67 Cash Advance Rate 0.03 0.16 0.02 159.92 Purchases \$166.86 \$318.86 \$192.91 86.50 Purchase Rate 0.74 1.29 1.81 40.62 Debits \$178.14 \$324.57 \$199.55 89.27 Debit Rate 0.77 1.31 1.84 41.99 Dollars in Segment 4.63 14.34 10.63% 43.53 Rate in Segment 3.32 9.64 11.89% 27.95

Detailed Description Paragraph Table (12):

TABLE 12 Segment Top 5% Scores Account Statistics: 154786 accounts (3.10 percent) Population Relative Category Mean Value Std Deviation Mean Score Cash Advances \$9.73 \$51.21 \$7.27 133.79 Cash Advance Rate 0.02 0.13 0.02 125.62 Purchases \$391.54 \$693.00 \$642.06 60.98 Purchase Rate 2.76 4.11 7.51 36.77 Debits \$401.27 \$702.25 \$649.34 61.80 Debit Rate 2.79 4.12 7.53 37.00 Dollars in Segment 1.24 8.14 1.55% 80.03 Rate in Segment 0.99 6.70 1.79% 55.04

Detailed Description Paragraph Table (13):

TABLE 13 Target population specification ID associated with promo- Segment Customer Selection tional offer ID target count Criteria Filter Criteria 1 122 75,000 Predicted Average Trans- Spending action in Seg- in Seg- ment >\$50 ment 1 143 Top 10% Dot Prod- Total Spending uct in Seg- ment >\$100 2 12 and 55 87,000 Predicted None Spending in this Segment 12 and 55

CLAIMS:

1. A method of predicting financial behavior of consumers, comprising: obtaining a set of input transactions for a plurality of consumers with respect to a plurality of merchants; defining at least one merchant segment, each merchant being associated with at least one of the defined merchant segments; and for at least one consumer, applying the input transactions of the consumer by a computer to each of at least one merchant segment predictive model, each merchant segment predictive model defining for a merchant segment a prediction function between input transactions in a past time interval and financial behavior in a subsequent time interval, to produce for each consumer a predicted behavior in each of at least a subset of the merchant segments.
2. The method of claim 1, wherein the predicted behavior comprises a likelihood of positive response to an offer.
3. The method of claim 1, wherein the predicted behavior comprises a spending level with respect to a merchant.
4. The method of claim 1, further comprising: generating a consumer vector for each of at least

a subset of the consumers; generating a merchant vector for each of at least a subset of the merchants; wherein defining at least one merchant segment comprises performing supervised segmentation on the merchant vectors.

6. The method of claim 4, wherein performing supervised segmentation comprises: initializing a set of segment vectors; accepting at least one segment label for at least one of the merchants; and for each of at least a subset of the labeled merchants: selecting at least one segment vector for a merchant having a merchant vector; determining whether the selected segment vector matches the segment label for the merchant; and responsive to the determination, adjusting zero or more of the segment vectors.

7. The method of claim 6, wherein selecting at least one segment vector for a merchant comprises selecting a segment vector that is closest to the merchant vector corresponding to the merchant.

8. The method of claim 6, wherein selecting at least one segment vector for a merchant comprises selecting at least one segment vector having a tolerance range that includes the value of the merchant vector.

9. The method of claim 1, further comprising: training a predictive model using the predicted behavior of a plurality of consumers in at least a subset of the merchant segments, and additional observed behavior for the plurality of consumers with regard to a target segment not included in the subset of merchant segments; and for at least one target consumer: providing as input to the trained predictive model predicted behavior of the target consumer in at least a subset of the merchant segments; and obtaining from the trained predictive model a predicted behavior of the target consumer with respect to the target segment.

10. The method of claim 1, further comprising: for at least one consumer, associating the consumer with the merchant segment for which the consumer had the highest predicting spending relative to other merchant segments.

11. The method of claim 1, further comprising: generating a consumer vector for each of at least a subset of the consumers; generating a merchant vector for each of at least a subset of the merchants; for at least one merchant segment, determining a segment vector as a summary vector of merchant vectors of merchants associated with the segment; and for at least one consumer, associating the consumer with the merchant segment having the greatest dot product between the segment vector of the segment and a consumer vector of the consumer.

12. The method of claim 1, further comprising: for at least one merchant segment: ranking the consumers by their predicted spending in the merchant segment; and determining for at least one consumer a percentile ranking in the merchant segment; and for each consumer: determining the merchant segment in which the consumer's percentile ranking is the highest, to uniquely associate each consumer with one merchant segment; and for at least one merchant segment, determining summary transaction statistics for the consumers uniquely associated with the merchant segment.

13. The method of claim 1, further comprising: for at least one merchant segment: ranking the consumers by their predicted spending in the merchant segment; determining for at least one consumer a percentile ranking in the merchant segment; selecting as a population, the consumers having a percentile ranking in excess of predetermined percentile threshold; and determining summary transaction statistics for selected population of consumers.

22. The method of claim 1, further comprising: generating a consumer vector for each of at least a subset of the consumers; generating a merchant vector for each of at least a subset of the merchants; determining for at least one merchant name in the transaction data a merchant vector; clustering the merchant vectors to form a plurality of merchant segments, wherein at least one merchant vector is associated with one and only one merchant segment; and for at least one merchant segment, determining from the transactions of consumers at the associated merchants of the merchant, statistical measures of consumer transactions in the segment.

23. The method of claim 1, further comprising: selecting a plurality of consumers associated with at least one merchant segment, the selected plurality selected according to their predicted spending in the merchant segment; and providing promotional offers to the selected

plurality of consumers.

24. The method of claim 1, further comprising: training at least one of the merchant segment predictive models to predict spending in a predicted time period based upon transaction statistics of the consumer's transactions in a past time period.

25. The method of claim 24, wherein the transaction statistics comprises variables describing the recency of the consumer's transactions in one or more merchant segments, the frequency of the consumer's transactions in one or more merchant segments, and the amount of the consumer's transactions in one or more merchant segments.

26. A system for predicting financial behavior of consumers, comprising: a database for storing a set of input transactions for a plurality of consumers with respect to a plurality of merchants; at least one merchant segment, each merchant being associated with at least one of the defined merchant segments; at least one merchant segment predictive model, for defining for a merchant segment a prediction function between input transactions in a past time interval and financial behavior in a subsequent time interval, to produce for each consumer a predicted behavior in each of at least a subset of the merchant segments.

27. The system of claim 26, wherein the predicted behavior comprises a likelihood of positive response to an offer.

28. The system of claim 26, wherein the predicted behavior comprises a spending level with respect to a merchant.

30. The system of claim 29, wherein the at least one merchant segment predictive model applies a learning vector quantization algorithm to the merchant vectors.

31. The system of claim 26, wherein: the merchant vector build module determines, for at least one merchant segment, a segment vector as a summary vector of merchant vectors of merchants associated with the segment; and the consumer vector build module associates at least one consumer with the merchant segment having the greatest dot product between the segment vector of the segment and a consumer vector of the consumer.

32. A computer-readable medium comprising computer-readable code for predicting financial behavior of consumers, the computer-readable medium comprising: computer-readable code adapted to obtain a set of input transactions for a plurality of consumers with respect to a plurality of merchants; computer-readable code adapted to define at least one merchant segment, each merchant being associated with at least one of the defined merchant segments; and computer-readable code adapted to, for at least one consumer, apply the input transactions of the consumer to each of at least one merchant segment predictive model, each merchant segment predictive model defining for a merchant segment a prediction function between input transactions in a past time interval and financial behavior in a subsequent time interval, to produce for each consumer a predicted behavior in each of at least a subset of the merchant segments.

33. The computer-readable medium of claim 32, wherein the predicted behavior comprises a likelihood of positive response to an offer.

34. The computer-readable medium of claim 32, wherein the predicted behavior comprises a spending level with respect to a merchant.

35. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to generate a consumer vector for each of at least a subset of the consumers; computer-readable code adapted to generate a merchant vector for each of at least a subset of the merchants; wherein the computer-readable code adapted to define at least one merchant segment comprises computer-readable code adapted to perform supervised segmentation on the merchant vectors.

37. The computer-readable medium of claim 35, wherein the computer-readable code adapted to performing supervised segmentation comprises: computer-readable code adapted to initialize a set of segment vectors; computer-readable code adapted to accept at least one segment label for at least one of the merchants; and computer-readable code adapted to, for each of at least a

subset of the labeled merchants: select at least one segment vector for a merchant having a merchant vector; determine whether the selected segment vector matches the segment label for the merchant; and responsive to the determination, adjust zero or more of the segment vectors.

38. The computer-readable medium of claim 37, wherein the computer-readable code adapted to select at least one segment vector for a merchant comprises computer-readable code adapted to select a segment vector that is closest to the merchant vector corresponding to the merchant.

39. The computer-readable medium of claim 37, wherein the computer-readable code adapted to select at least one segment vector for a merchant comprises computer-readable code adapted to select at least one segment vector having a tolerance range that includes the value of the merchant vector.

40. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to train a predictive model using the predicted behavior of a plurality of consumers in at least a subset of the merchant segments, and additional observed behavior for the plurality of consumers with regard to a target segment not included in the subset of merchant segments; and computer-readable code adapted to, for at least one target consumer: provide as input to the trained predictive model predicted behavior of the target consumer in at least a subset of the merchant segments; and obtain from the trained predictive model a predicted behavior of the target consumer with respect to the target segment.

41. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to, for at least one consumer, associate the consumer with the merchant segment for which the consumer had the highest predicting spending relative to other merchant segments.

42. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to generate a consumer vector for each of at least a subset of the consumers; computer-readable code adapted to generate a merchant vector for each of at least a subset of the merchants; computer-readable code adapted to, for at least one merchant segment, determine a segment vector as a summary vector of merchant vectors of merchants associated with the segment; and computer-readable code adapted to, for at least one consumer, associate the consumer with the merchant segment having the greatest dot product between the segment vector of the segment and a consumer vector of the consumer.

43. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to, for at least one merchant segment: rank the consumers by their predicted spending in the merchant segment; and determine for at least one consumer a percentile ranking in the merchant segment; and computer-readable code adapted to, for each consumer: determine the merchant segment in which the consumer's percentile ranking is the highest, to uniquely associate each consumer with one merchant segment; and for at least one merchant segment, determine summary transaction statistics for the consumers uniquely associated with the merchant segment.

44. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to, for at least one merchant segment: rank the consumers by their predicted spending in the merchant segment; determine for at least one consumer a percentile ranking in the merchant segment; select as a population, the consumers having a percentile ranking in excess of predetermined percentile threshold; and determine summary transaction statistics for selected population of consumers.

53. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to generate a consumer vector for each of at least a subset of the consumers; computer-readable code adapted to generate a merchant vector for each of at least a subset of the merchants; computer-readable code adapted to determine for at least one merchant name in the transaction data a merchant vector; computer-readable code adapted to cluster the merchant vectors to form a plurality of merchant segments, wherein at least one merchant vector is associated with one and only one merchant segment; and computer-readable code adapted to, for at least one merchant segment, determine from the transactions of consumers at the associated merchants of the merchant, statistical measures of consumer transactions in the segment.

54. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to select a plurality of consumers associated with at least one merchant segment, the

selected plurality selected according to their predicted spending in the merchant segment; and computer-readable code adapted to provide promotional offers to the selected plurality of consumers.

55. The computer-readable medium of claim 32, further comprising: computer-readable code adapted to train at least one of the merchant segment predictive models to predict spending in a predicted time period based upon transaction statistics of the consumer's transactions in a past time period.

56. The computer-readable medium of claim 55, wherein the transaction statistics comprises variables describing the recency of the consumer's transactions in one or more merchant segments, the frequency of the consumer's transactions in one or more merchant segments, and the amount of the consumer's transactions in one or more merchant segments.

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)